S. HRG. 100-947

COMPUTER NETWORKS AND HIGH PERFORMANCE COMPUTING

HEARING

BEFORE THE

SUBCOMMITTEE ON SCIENCE, TECHNOLOGY, AND SPACE

OF THE

COMMITTEE ON COMMERCE, SCIENCE, AND TRANSPORTATION
UNITED STATES SENATE

ONE HUNDREDTH CONGRESS

SECOND SESSION

ON

COMPUTER NETWORKS AND HIGH PERFORMANCE COMPUTING

AUGUST 11, 1988

Printed for the use of the
Committee on Commerce, Science, and Transportation

U.S. GOVERNMENT PRINTING OFFICE
91-313 O WASHINGTON : 1988

For sale by the Superintendent of Documents, Congressional Sales Office
U.S. Government Printing Office, Washington, DC 20402

S261-19

COMMITTEE ON COMMERCE, SCIENCE, AND TRANSPORTATION

ERNEST F. HOLLINGS, South Carolina, *Chairman*

DANIEL K. INOUYE, Hawaii
WENDELL H. FORD, Kentucky
DONALD W. RIEGLE, JR., Michigan
J. JAMES EXON, Nebraska
ALBERT GORE, JR., Tennessee
JOHN D. ROCKEFELLER IV, West Virginia
LLOYD BENTSEN, Texas
JOHN F. KERRY, Massachusetts
JOHN B. BREAUX, Louisiana
BROCK ADAMS, Washington

JOHN C. DANFORTH, Missouri
BOB PACKWOOD, Oregon
NANCY LANDON KASSEBAUM, Kansas
LARRY PRESSLER, South Dakota
TED STEVENS, Alaska
ROBERT W. KASTEN, JR., Wisconsin
PAUL S. TRIBLE, JR., Virginia
PETE WILSON, California
JOHN McCAIN, Arizona

RALPH B. EVERETT, *Chief Counsel and Staff Director*
WALTER B. MCCORMICK, JR., *Minority Chief Counsel and Staff Director*

SUBCOMMITTEE ON SCIENCE, TECHNOLOGY, AND SPACE

DONALD W. RIEGLE, JR., Michigan, *Chairman*

ALBERT GORE, JR., Tennessee
JOHN D. ROCKEFELLER IV, West Virginia
LLOYD BENTSEN, Texas
JOHN F. KERRY, Massachusetts
BROCK ADAMS, Washington

LARRY PRESSLER, South Dakota
NANCY LANDON KASSEBAUM, Kansas
PAUL S. TRIBLE, JR., Virginia
PETE WILSON, California

CONTENTS

Page
Opening statement by Senator Gore ... 1
Opening statement by Senator Riegle ... 4
Opening statement by the Chairman ... 7

LIST OF WITNESSES

Connolly, John, professor of physics, University of Kentucky ... 80
  Prepared statement ... 83
  Questions of the Chairman and the answers ... 88
Dickinson, Dr. Robert E., deputy head, interdisciplinary climate systems section, National Center for Atmospheric Research ... 46
  Prepared statement ... 47
Huray, Dr. Paul G., vice president for research, University of South Carolina ... 8
  Prepared statement ... 13
Kahn, Dr. Robert E., president, Corporation for National Research Initiatives ... 35
  Prepared statement ... 39
  Questions of the Chairman and the answers ... 42
King, Dr. Kenneth M., president, Educom ... 60
  Prepared statement ... 64
  Questions of the Chairman and the answers ... 69
Kleinrock, Dr. Leonard, professor of computer science, University of California ... 18
  Prepared statement ... 20
  Questions of the Chairman and the answers ... 23
Kutler, Dr. Paul, Acting Director, Aerodynamics Division, Office of Aeronautics and Space Technology, NASA ... 52
  Prepared statement ... 56
McAdams, Dr. Alan, associate professor, Johnson Graduate School of Management, Cornell University ... 98
  Prepared statement ... 101
  Questions of the Chairman and the answers ... 117
Wulf, Dr. William, Assistant Director, Directorate for Computer and Information Science and Engineering, NSF ... 90
  Prepared statement ... 93

ADDITIONAL ARTICLES, LETTERS, AND STATEMENTS

Rona, Dr. Thomas P., Deputy Director, Office of Science and Technology Policy, statement ... 131
Toward A National Research Network, article ... 132
U.S. Supercomputer Vulnerability, article ... 136

COMPUTER NETWORKS AND HIGH PERFORMANCE COMPUTING

THURSDAY, AUGUST 11, 1988

U.S. Senate,
Subcommittee on Science, Technology, and Space,
Committee on Commerce, Science, and Transportation,
Washington, DC.

The subcommittee met at 9:55 a.m., in room SR-253, Russell Senate Office Building, Hon. Albert Gore, Jr., presiding.

Staff members assigned to this hearing: Pat Windham and Mike Nelson, professional staff members; and Mark Farrell, minority professional staff member.

OPENING STATEMENT BY SENATOR GORE

Senator Gore. The subcommittee will come to order.

I want to apologize to our witnesses and guests for the delay in beginning. As may have been explained to you, the Senate was scheduled to vote at precisely the same time as this hearing was scheduled to begin and I felt the way to have the least delay would be to make that vote and quickly come over here and start the hearing.

As sometimes occurs, there was unanimous consent to have an extra five minutes of debate and then an extra ten minutes of debate, so the vote just began a few moments ago. I came right over here. In any event, there was no discourtesy intended. Quite the contrary, we are trying to make it as short a delay as possible.

I am delighted to welcome all of the witnesses here today. Because of our late start, I am going to put my own prepared statement into the record and just summarize portions of it very briefly.

I believe very strongly that ensuring America's world leadership in advanced computer technology may well be the most important economic and technological challenge of the 21st century. Where we once led the world in every technology from automobiles to integrated steel, we now face much stiffer global competition.

High performance computing represents one of our best hopes to improve American competitiveness by restoring growth in productivity.

Our country desperately needs to maintain basic manufacturing. But basic industry demands much more than it used to. Many of our factories have fallen behind the rest of the world because much of our technology is out of date and inflexible. Powerful supercomputers can change the way America does business.

Properly harnessed and directed, massive computing power can do much more: from predicting global climate changes and designing efficient automobile engines to testing new molecules for medical value. If we can rise to these grand challenges of supercomputing, America will excel in commerce and technology and science well into the next century.

Two years ago I sponsored the Supercomputer Network Study Act to explore the potential of fiber optic networks, high-capacity fiber optics networks, to link supercomputers into one system. I introduced the bill on the 30th anniversary of the interstate highway system because I believe high-capacity fiber optic networks will, in fact, be the information superhighways of tomorrow.

The legislation was attached to the 1986 National Science Foundation authorization and the study of fiber optics evolved into the Office of Science and Technology Policy's wide-ranging report on high performance computing.

When I introduced my legislation, I envisioned a national computer network linking academic researchers and industry, clustering research centers and businesses around network interchanges, and using the Nation's vast data banks as the raw material for increasing industrial productivity and creating new products. Today, we have the ability to create an information Interstate Highway for the 1990's, to transport information, rather than goods, at nearly the speed of light to be processed and adapted for industrial productivity and scientific research.

I think it is abundantly obvious for a number of reasons that our Nation will benefit greatly if we can put in place a high-capacity network linking the Nation's supercomputing centers.

Today and in other hearings, I want to look at what researchers need, ways to pay for an improved National Research Network, whether or not government agencies are working together well to develop the network that all can use, and where we are in relation to other nations' networks.

There are already over 100 networks in this country, but the coordination among them is limited. Network users attest to the difficulty of shifting among networks to get to the right data bank or the right supercomputer or the colleague they wish to work with on a particular project. Despite these problems, networks already allow researchers to share ideas, data, computer programming and documents.

In superconductor research, for example, where discoveries are made every week and sometimes several times a week, computer networks and fax machines have now virtually replaced journal articles as the primary means of sharing results. Researchers without networks or telecopiers are lost in the fax gap.

According to the Office of Science and Technology Policy report, and I quote, "National networks have a low capacity, are overloaded and fail to interoperate successfully." The report goes on to say that the performance of the underlying network technology has not kept pace with the increased demands.

Therefore, the networks which in the 1970's had significant impact in enabling collaboration are now barriers.

The OSTP report also notes that "the basic technologies for later generations (of national computer networks) are also being developed in the United States, but there have been no major efforts to apply them to address the needs." Once again, we are in danger of virtually inventing a technology only to see other nations develop and profit from it.

In computer technology the hardware is typically several years ahead of the software. I think it is safe to say that the software is typically many years ahead of the policy. Just as it takes time for programmers to design applications for new machines, it takes time for policymakers to address the problems that new technology brings.

We cannot afford to delay any longer. We must begin to anticipate problems with greater success even though we are just starting to realize the productivity gains the computer revolution can provide.

We must develop a national plan to take advantage of our computing technology and to cope with problems of computer security, software copyrights and many other problems.

Many of the witnesses today will discuss the national research network, NSFNET, which promises to enhance coordination among researchers and universities and to allow greater access to supercomputers. But many problems remain, including slow transmission speed, lack of access, and budgetary constraints. In addition, we must decide who will manage the network and make policy decisions about the entire system. We must protect the system against viruses and breaches of security like the recent invasion of Los Alamos National Lab data banks by East German teenagers.

Currently, the five national supercomputer centers funded by NSF and the National Center for Atmospheric Research are linked, forming the backbone of NSFNET. This July, managers increased data transmission from 56 kilobits per second to 1.5 megabits per second, the equivalent of 50 pages of single-spaced typewritten text a second. By 1989, more than 200 universities will have access to the network. And a year later, NSF expects to boost data transmission speed to 45 megabits, or the equivalent of an entire Sunday newspaper every second.

In addition to the $500 million a year that Federal agencies currently spend on high-performance computing, the OSTP recommended that the Federal Government spend an extra $1.7 billion over 5 years. Of this increase, $390 million over 5 years would be for the development of the National Research Network. We currently spend $50 million a year. Furthermore, NSFNET has a 5-year $14 million grant from NSF and has received $25 to $30 million in goods and services from IBM and MCI. OSTP's funding recommendations are aimed at developing a three gigabits per second network within 15 years.

However, the National Academy of Science report on the national research network points out that the first two stages of OSTP's three stage network recommendation will merely enable us to catch up to state-of-the-art technology—and that the third stage of three gigabits per second needs more research now if we are ever to achieve that goal. The NAS report pointed out other concerns within the network plan. Where the OSTP report emphasized speed, the NAS report also underscored the need for proper management, adequate coordination, and a more complete analysis of funding proposals.

These hearings will delve into these critical questions. Supercomputers will allow us to study possible pharmaceutical molecules, model astronomical phenomena, analyze computer chip performance, enhance oil recovery, refine three-dimensional medical imaging, and analyze fusion. Linked by vast networks—the infrastructure of the 21st century, supercomputers will also allow us to design efficient car engines, aircraft, and appliance; develop a hazardous waste emergency data base; forecast weather more accurately and much longer in advance. With high-speed computer networks, a surgeon in Nashville can send a CAT scan picture to a colleague at the Mayo Clinic and get a second opinion instantly. A researcher can now sit in an office and, with a few keystrokes, access a data base thousands of miles to produce a map of any neighborhood in the country, detailed enough to show individual houses. A particle physicist in California can check up on an experiment being run at Fermilab in Illinois without having to leave his office. A historian can track down a reference in seconds.

This morning, we can only imagine all that this powerful tool may one day accomplish. I hope that when this hearing is over we will all know a little more about how to make a national supercomputers network a tool to build America's future.

I want to thank the chairman of the Subcommittee, Senator Don Riegle. We served together in the House, and I want to thank him for allowing me to take this initiative and for his fine work in this area as well. And I recognize him at this time.

OPENING STATEMENT BY SENATOR RIEGLE

Senator RIEGLE. Thank you, Senator Gore.

Let me say I remember with great fondness our days together in the House and I really appreciate the fact we are able to serve together in the Senate.

I want to begin by commending your leadership in this area. As chairman of the Subcommittee on Science, Technology, and Space, I have encouraged various members of the Subcommittee with particular expertise and background to take the lead in areas of particular interest and concern to them. Senator Gore, your leadership in this area has been, I think, in the forefront of the entire Senate and so I appreciate very much the initiative and the hearing that has been put together today.

I think this is a very important hearing and I think that building this Committee record and building the policy foundation needed to address some of the issues that have already been raised is vital to the future of this country. We will be receiving today national trade data again which will show that the United States continues to lag far behind in the international trading system.

We have just passed a comprehensive trade bill with some real teeth in it to try to stop the trade cheating and get us more up to speed with respect to the competitive aspects of the world economy, but at the foundation lies our ability to adapt ourselves rapidly to technological change and to be able to put ourselves in a position to really meet world competition head on.

We know that countries like Japan, for example, in high technology applications have done an extraordinary job and they continue to press very hard.

We cannot afford any inefficiency in our system, whether it is insufficient training, capacity in our people or any lags or breakdowns in the process of scientific advancement. We need to cross-connect our advances so that we are able to move as a nation at the most rapid rate of speed.

These issues relate to our capacity to gather information and share it widely throughout our society, to move it at the greatest speed, and to get it to where it is needed so that it can make its way into our economic decisionmaking and our economic activities. So, addressing these issues is absolutely crucial to our future.

As in biotechnology and many other areas, Senator Gore, where you have been out front, your leadership here is very important.

I want to touch on a few other things in my opening statement as well.

I think in order to maximize the aggregate computing power that is available to the nation, efficient networks must be established to link computers which can permit the exchange of enormous amounts of data in very short periods of time.

Already some networks of this nature have been established. One of these is the NSFNET which has been headquartered in my own State of Michigan.

This system will link five NSF supercomputer centers with one communications system. This is the most advanced network of its kind in the world and permits scientists from all regions of our country to share data and to work on common problems.

Clearly the potential for scientific advancement is very great and NSF should be commended for implementing this program.

It is also, I think, significant to note that private industry has made a large contribution to the success of this program, most notably IBM and MCI, in addition to a major financial commitment from the State of Michigan itself.

Despite that example of a success, however, there are aspects in the current networks which limit their growth. One such area involves the ability of the telecommunications system to accommodate large amounts of data. It is limited in many respects to the capacity of copper wire phone lines as opposed to the fiber optics employed in the NSFNET system.

In order to fully realize the potential of the supercomputer networks and to reap fully the benefits of these networks, the capacity to transmit data has to be greatly expanded. This will be expensive. But I think it will be cost-effective and it is an investment in our future that is unavoidable if we are going to create these systems and achieve their full potential.

So the hearing today will examine the current state of computer networking in this country, the immediate prospects for improving these networks, the cost of these improvements and where other nations stand in relation to the United States in this technology.

We cannot lose sight of the fact that the key to continued scientific, technical and industrial advancement will be the ability to tie our computing systems together. The French, for example, have linked more than 4 million small computers using hardwire technology. The Japanese are investing large amounts in fiber optic networks. So, we as a nation must understand and face and respond to the challenge that we are seeing by our foreign competitors in the world economic system.

Finally, I just want to make one other comment. The State of Michigan which I represent, is an advanced manufacturing State, which has felt as much as any State the shocks caused by changes in the world economic system. And we have been adjusting to those changes, I would say to the witnesses, Senator Gore, and others present today, I think quite well.

We have managed to bring our unemployment rate in Michigan down very close to the national average. We still have pockets of difficulty but we are now applying the highest levels of technology to manufacturing, and are managing the transition to high value-added production, which over the years has been critical to the economic success of the country as a whole. We are actually building the middle class and giving people the ability to have a broadly shared middle class standard of living.

In order to continue to advance and at the same time reduce the $170 billion trade deficit, we must consider how our scientific people can improve our capacity to use the breakthroughs in science to enhance our ability to produce end products. This is now something that has to be part of every decision we make.

We cannot afford to leave out the considerations of how we can apply ever-increasing and more sophisticated amounts of technology to the capacity of each individual worker.

Obviously the technology has to be, as we say, worker friendly because it has to be something that really enables the individual worker to produce more and produce it more quickly. But all the way up and down the line, from the scientific inquiry side of it through advanced scientific work right down finally to the plant floor and the job site, we have to be able to link our capacity to think and work and apply information and new ways of doing things in order to be able to excel and to be able to meet and exceed the gains that we see our foreign competitors making.

So that is really the challenge to this country. I would say that in the last few years there has not been sufficient emphasis on this within the broad strategy of our national government. This Committee and Senator Gore and others have tried to lead that effort. I would hope in the immediate future—we are about to have a change in national administration—that we would see an opportunity to take the kind of work and findings that will be developed today and lay out a policy strategy for the future that can really begin to multiply the gains that will come into the day-to-day lives of working citizens and families across this country.

Senator GORE. I thank you very much again for your leadership in this area.

Thank you for all the work you have done. I certainly appreciate the opportunity to have this hearing today.

I have a statement that the Chairman of the Committee, Senator Hollings, would like to have included in the record at this point. [The statement follows:]

OPENING STATEMENT BY THE CHAIRMAN

Today, we live in the Information Age. There are entire industries, like insurance and banking, that deal not in physical commodities, but in information—information which computers help us manage and process.

Today, the Science, Technology, and Space Subcommittee will examine one aspect of the computer revolution. The subcommittee will explore the latest developments in computer networking. Just as people need telephones, computers need networks to communicate. Since the first national computer network started less than twenty years ago, networking has gone from being a curiosity to a billion-dollar business. Every time you use an automated teller machine or order an airline ticket, you're relying on a computer network.

Today's hearing will focus on more advanced computer networks. In the last few years, scientists have developed networks that allow computers to exchange a million of bits of data per second. That's equivalent to fifty pages of single-spaced text in a single second. There are experimental networks that can transmit the entire contents of the Encyclopedia Britannica in five seconds. Such experimental systems will become ordinary in a few years.

By the mid-1990's, many home computers will be linked to these high-speed networks. In the future, your computer may help you plan your day before you leave for work. You might turn on your computer in the morning and get the weather map and the forecast for Washington, D.C., or any other city in the world. Then you could call up a map of traffic conditions so you can beat the traffic on the way to work. Then you might decide that you need to go shopping for a new jacket after work. Your computer could search a file of clothing store ads and find you the best deal on the jacket you want. At the same time, you could access data bases to help you pick a restaurant for lunch, plan a vacation, or check you bank account. In just a few minutes, you might be able to save yourself an hour or two of hassle.

While this kind of networking technology is not yet available in the home, many scientists are using high-speed networks and data bases to make their work easier. Today a scientist without access to a network is like a Senator without access to a microphone. This technology is already allowing scientists around the country to share new data the day it is collected, rather than relying on the mail. Ten years ago a physicist would mail a colleague a stack of computer paper covered with numbers and then wait a week or two for a reply. Today, that same physicist can use his computer to send a three-dimensional plot of his data to his colleague across the country and they can talk on the phone while plotting the data in different ways until they've figured out what it's telling them. This instantaneous sharing of data makes it possible for experiments to be done in a few days, rather than a few weeks. And it enables scientists to work together even if they are thousands of miles apart. This is becoming more and more important because we are increasingly faced with scientific problems that require dozens of scientists to work closely together. Articles on high-energy physics experiments often will have forty co-authors who helped make the experiment possible. We will not be able to cope with the problem of greenhouse warming without the combined efforts of atmospheric physicists, chemists, oceanographers, supercomputer programmers, ecologists, and many others. Computer networks will let researchers do collaborative research without spending hours each month in airplanes and airports.

The committee will hear about the NSFNET and about how it allows scientists throughout the country to access the five NSF supercomputer centers. Now researchers at small colleges in South Carolina will have access to the same computing power available at places like Harvard or Caltech. Thanks to NSFNET, researchers at these schools will be able to use the most advanced computers in the world to develop methods for oil exploration, to design new wonder drugs, or to improve our ability to predict the weather. A small school can't afford to spend millions of dollars on a supercomputer, but for a few thousand dollars it can become part of NSFNET and take advantage of this new technology.

The NSFNET represent the first phase of an advanced national research network. Will this network be well-funded and well-coordinated or will it be unreliable, hard-to-use, and unable to meet the needs of researchers? The answer will depend upon whether the network is properly planned and adequately funded. The subcommittee is meeting today to get the information it needs to decide the best way to fund and build a national computer network. We need to take the time to examine the critical management and budgetary problems that must be tackled. For a national computer network is an important research tool, which unlike a telescope or a particle accelerator, will help scientists from all disciplines.

Without an adequate national computer network, American researchers, and indeed all Americans, will not be able to enjoy all of the benefits of the computer revolution. Other countries may reap those benefits instead. We must strive to take full advantage of this new, promising technology.

Senator GORE. Let's move right onto the first panel. I would like to call Dr. Paul Huray, Chairman of the Federal Coordinating Council for Science, Engineering and Technology Committee on Computer Research Applications and Leonard Kleinrock, Chairman of the National Research Network Review Committee of the National Research Council.

Dr. Huray is vice president for Research at the University of South Carolina now and, most important of all, a native of Tennessee. Dr. Kleinrock is professor of Computer Science at the University of California in Los Angeles.

I am delighted that you are both here and I respect the work you have done. Dr. Huray, why do we not begin with you.

STATEMENT OF DR. PAUL G. HURAY, CHAIRMAN, FEDERAL COORDINATING COUNCIL FOR SCIENCE, ENGINEERING, AND TECHNOLOGY, COMMITTEE ON COMPUTER RESEARCH AND APPLICATIONS AND VICE PRESIDENT FOR RESEARCH, UNIVERSITY OF SOUTH CAROLINA

Dr. HURAY. Thank you, Senator Gore. It is a pleasure and honor to appear here today before you and Mr. Riegle and members of the staff to report on the work of the Committee on Computer Research and Applications. I have prepared remarks and I would like to merely summarize them.

Senator GORE. Without objection, we will put the prepared statements of all the witnesses today into the record and if you wish to summarize, that would be fine.

Dr. HURAY. I plan to concentrate today on the report produced for the Congress as a result of your legislation, 99-383, entitled "A Research and Development Strategy for High Performance Computing"[1] which has been delivered to your staffs.

I want to begin by recognizing the important work of this Committee and the staff with whom we have worked well, especially the work of the chairman of the staff, Pat Windham. We found strong cooperation between the Congressional and the Executive Office and I also want to recognize the strong cooperation that we have had among the agency participants in this report who are listed on the last page. They are among the best computing people in the agencies and we have reached consensus in preparing this document.

I would like to concentrate primarily upon the two main points in the report dealing with balance and the participation of industry. Our report recommends four primary recommendations. They are categorized in the areas of high performance computer hardware or architectures, software, algorithms, scientific applications, basic research in human resources and the National Research Network which we very strongly support.

We think the balanced implementation will give us the maximum impact as a nation and the participation of industry will leverage Federal dollars and make the results very realistic. We think by coming forward with a balanced implementation we can continue leadership in high performance computing and increase productivity of the science and technology community.

---

[1] Printed by the Executive Office of Science and Technology Policy, Nov. 20, 1987, 29 pages.

In short, we think this is a "good news strategy: We have the lead and we intend to keep it."

I would like to remark briefly on each of the areas:

In the area of *architectures*, the OTA a few years ago produced a report which showed that over the last 30 years the performance of our highest performance computers had been increasing by an order of magnitude every five years, (roughly). This improvement has been a result primarily of new ways of packaging and cooling chips, making them more dense. We have been limited primarily by the production of pure materials and step permoters for the lithography process, but people are beginning to realize that we are reaching the physical limits of making smaller and smaller components.

We believe that we can sustain, however, historical growth patterns with the new architecture provided by massive parallelism. Our report speaks strongly towards supporting that continued research.

Senator GORE. If I could interject, the physical limits you refer to are the speed of light and the distance between molecules?

Dr. HURAY. Yes, sir. The numbers of atoms involved in some of these chips are beginning to be finite and so we are reaching a limit that way. It means that we can make these components only so close to one another and the speed of light will then determine the speed at which the clock operates. The faster clocks will not continue to be the only answer but computers working in parallel with one another will continue to produce higher speeds.

As you pointed out in your remarks, Senator Gore, we lead the world in the production of software but, unfortunately, much of it in the United States is written in garages and university offices. It typically deals with small isolated problems. It typically does not conform to a larger standard. It is not portable in many cases from one computer to another. It does not scale from a PC to a mainframe and it does not lend itself well, in many cases, to parallelization. That is a new word I invented.

As you correctly said, software is leading our hardware development. We believe the National Research Network will help alleviate some of these problems and promote the sharing of software among researchers. We think it is also important, however, to focus upon some of the large problems in science and technology with software groups. In accordance with the term coined by Ken Wilson, "the grand challenges," we think it is important to focus on generic computing tools and algorithms that any research scientist can use, whether he has a background in computer science or not.

I note in my testimony some important areas in which software is focusing these days. New aircraft and automobile design is being performed with some of these activities and you are going to hear about this in later testimony. There are efforts in mineral recoveries, weather forecasting, materials science, and development of specific pharmaceuticals. We believe the focus on these activities will have a very large leveraging effect, well in excess of that of the high performance computer hardware, which is about a $1 billion a year industry.

In terms of basic research and human resources, we think that computational problems will always continue to occur just beyond our capability and that basic research will continue to provide new ideas for exploitation. So we feel it is important to continue to focus on the basic research in concert with a high performance network.

We think it is important to recognize the need to fuel this activity with manpower, particularly in light of changing demographics in the United States. We feel it is important to give access to many universities where some of our finest students are located and especially to give access to institutions where minorities are participating in research since they are becoming a larger part of our university population.

Now, regarding networking, I made an analogy to the early days of the telephone system. Apparently several companies competed for service in the early days. Typically a house would have several phone sets; one for each company and each of the phone sets would have a different number. Typically they did not interconnect, there were no common directories and people could not easily communicate with one another.

The situation in research networking is much like that today. The interconnect among the networks is poor and protocols are often different. Sometimes researchers have five or six identification numbers. We visited departments of science on university campuses in which four or five connections would be made to different networks and often those researchers would not be aware of one another.

There is usually no common directory for those networks and it is like the Moscow phone system; you have to ask, "what is your telephone number?"

The National Academy Report speaks to this and I will not comment further.

One of the ironies of producing our report was I was in Tennessee working with many people here in Washington, we were trying to send the report back and forth over BITNET and we found often it came out jumbled on the other end because it was a large report. So the best way for us to communicate was via Federal Express.

Senator GORE. That was the network used by the Network Committee?

Dr. HURAY. That's correct. We do currently have several dozen existing agency networks which we think can be used to build an internet and begin the first phase of the National Research Network. There are also several national commercial fiberoptic systems which we think should participate in this growth process.

To achieve this internet, we have produced a Committee (a subset of the FCCSET Committee), called the FRICC, Federal Research Internet Coordinating Committee, which has begun the first stages of implementing the National Research Network. They have accomplished this by memoranda of understanding between NSF's NSFNET; DARPA's ARPANET; and NASA's NASNET. We expect to expand to other agencies over time. And as you know, a few weeks ago there was an announcement in Congress that the backbone for that network had been raised to 1.5 megabits per second. Typically, however, users at individual universities are still limited to 1.2 or 9.6 kilobits per second.

We see a three-tiered hierarchy being developed, sort of a network hierarchy: There is a national backbone, many regional networks and then local institutional networks on campuses and national laboratories. We think that this association should become the National Research Network over time. I have described the three phases of the network heirarchy in my comments and I am not going to repeat them.

We think there are several opportunities for a national network that are above and beyond those that we originally recognized and that we thought would be the primary justification. Originally we believed access to the supercomputer centers by many students across the country would be the primary benefit of the network. We still believe that is one of the primary benefits. However, we now believe equally important to be the process of collaboration among scientists. We think that one day scientists will communicate among one another, collaborate and exchange software in much the same way that they pick up a telephone and talk to one another today.

We have also learned about a tremendously exciting group of innovators who we believe will increase our productivity in the future. There are people today performing research in distributed hypercubes around the country in which entire computers act only as nodes in a massive parallelism. We see the potential for those growing to the point where one day we might wheel around computer power much in the same way we wheel electrical power on electrical networks today.

Senator GORE. Now, a hypercube, excuse me for interrupting, that word is a way of describing massive parallel computers where computer processors are in parallel and able to relate not just on one plane with the other processors but in several different dimensions?

Dr. HURAY. Yes, sir. It is one of the techniques being used for massive parallelism. There are several other techniques also being explored.

There are also researchers working on artificial intelligence developing mathematical constructs which we call "knowbots." Knowbots are constructs that are launched to find information in databases around the country, sometimes those written in different protocols and often having disparate information. If the networks are semi-intelligent we might one day be capable of retrieving information from several geographically distributed databases around the country and having a mathematical construct return with its information; much in the same way a robot does that today in a warehouse.

Senator GORE. That is k-n-o-w bots?

Dr. HURAY. Yes, sir. We have been familiarly calling those "Obewan k-n-o-w bots" because we want to provide a positive connotation for this process.

Of course this effort will not be free and we have provided a realistic budget estimation in our report. We think that the additional investment by the Federal Government must be matched in several ways by industry and states. We have taken not just a more-of-everything approach. We think that the economy of scale will be important and, in fact, we believe we have already demonstrated that in achieving Phase I of the National Research Network, within existing budgets.

However, we think that the real economic payback will not be just in a capability of researchers to communicate with one another or in sales of computers but in the tremendous leverage found in other technologies. Since the report was issued, we have begun budget cross cut analysis with the assistance of OMB, and Dr. William Graham has sent a memorandum to agency heads giving his support to this activity. We expect to come forward in the fiscal year 1990 budget with the first phase of the coordinated budget recommendation.

Of course, we recognize that such a strategy is difficult to achieve. However, it has many strengths, one of which is that the network is geographically distributed. It is built from existing agency networks and does not require decisions about big science versus little science. Because we have recommended budgets which are distributed among agencies, we believe those budgets can be stable over time and they can also meet agency priorities.

This implementation will, of course, require the continued working support of both the Congress and the Executive Office. We think the paybacks to the American taxpayer will be enormous. Thank you, Senator Gore.

[The statement follows:]

STATEMENT OF PAUL G. HURAY, SENIOR VICE PRESIDENT FOR RESEARCH, UNIVERSITY OF SOUTH CAROLINA

Chairman Gore and members of the Committee, it is an honor to appear before you today to discuss the needs and opportunities of a National Research Network. I shall confine my remarks to the work of the multiagency FCCSET Committee on Computer Research and Applications under the aegis of the Office of Science and Technology Policy and in particular to the report entitled, A Research and Development Strategy for High Performance Computing, which was delivered to the Congress on November 20, 1987 in response to PL 99-383.

Before I begin those remarks, I want to state the strong support of the FCCSET Committee for a National Research Network and recognize the leadership that the Senate Committee on Commerce, Science and Transportation and it's staff has provided to lend impetus to these studies. A continued working partnership between Congress and the Executive branch is essential to accomplishing our common goal of continued leadership in science and technology.

I also feel it important to apprise you of the strong cooperation that has been shown by the various agencies through their representatives (who are listed on the last page of the November 20, 1987 report). Because there is no single private corporation or Federal agency with a mission for serving the computing needs of the entire U.S. research community, I feel that the FCCSET Committee should continue to play a strong coordinating role in the creation of a National Research Network under the direction of the Office of Science and Technology Policy. Finally, I would mention the strong interagency coordination which has taken place under the leadership of the FCCSET and Dr. William R. Graham as chairman.

BALANCE

One of the most important aspects of the November 20 report was the determination to place the National Research Network into a four part set of recommendations:

- High performance computer hardware or architectures,
- Software, algorithms, and scientific applications,
- Basic Research and Human Resources, and
- Networking.

I feel that a balanced implementation of the strategy along all four parts will ensure the maximum impact upon the research infrastructure. It will provide a network to best take advantage of the nation's investment in computing, allow maximum access to advanced hardware and data bases, transmit the most useful information, maintain growth and continually rejuvenate the computing capacity of the nation. Furthermore, the involvement of industry in the implementation will ensure the efficient use and high leveraging of Federal dollars. I maintain that the integrated strategy will allow us to continue leadership in computing technology, have a positive effect upon our scientific and technological productivity, and assist a powerful cultural change in the United States research and development infrastructure for the future. In that sense, this is a "good news strategy: we have the lead and we intend to keep it."

ARCHITECTURES

A recent Office of Technology Assessment study has shown that the computer industry has continued to advance the state of the art in high performance computing by roughly an order-of-magnitude improvement each 5 years through aggressive use of packaging and cooling thus producing denser, more compact chip components. The ability to fabricate the components has been limited primarily by our ability to produce pure materials and precise stepper motors for photolithography. But we are now beginning to approach some physical limits of nature. The distance between chip elements is beginning to be limited by the wavelength of the light that is used in the lithography process, the width of the elements is measured in small numbers of atoms, and the overall performance is measured by the time required for light to travel from one point to another. It now appears that if we are to sustain the rapid performance increases which have been made in computer hardware in the past 30 years we can best do that through applications of new architectures which are in our research laboratories and under development in our industry today. With the new paradigm of architectures based upon massive parallelism, many individual computers working on aspects of the same problem simultaneously, the performance of state-of-the-art computers can be made to run even faster. We are projecting that today's supercomputers which operate at tens of billions of operations per second will, by the turn of the century, operate at trillions of operations per second (the so-called TERAOP computer) through the use of parallelism.

SOFTWARE

Even though the United States leads the world in software development, much of it is still written in garages or university offices. Today's software typically addresses small isolated problems, does not conform to a larger standard, is not portable from computer to computer, does not scale from a P.C. to a main frame, and does not lend itself naturally to parallelization. Our report concluded that software is lagging our hardware development. The National Research Network will help alleviate some of these problems by promoting the sharing of software among researchers but it is equally important for us to focus some of our efforts on particular computational missions (such as the grand challenges) and encourage the development of generic computing tools and algorithms which can be used easily by the average researcher. Using these tools which have an overlap with industry will have a long-range payoff to the American taxpayer. The phenomenal rate of change of our hardware has already produced a computational branch of science to complement experiment and theory. In some cases it allows us to carry out hypothetical experiments at velocities which could not be accomplished in a laboratory or to compress or expand time. Further advances in performance via special architectures and new software tools will allow us to apply computation to aircraft or automobile design, mineral recoveries, weather forecasting, materials development, and the design of specific pharmaceuticals that will leverage manyfold the investment we make today

BASIC RESEARCH AND HUMAN RESOURCES

Although the FCCSET committee recommended three phases of implementation for a National Research Network and specific approaches to hardware and software, there will always be computational problems just beyond our capability. Basic research will continuously provide new computational opportunities for exploitation. Individuals, small theory groups, and large consortia will provide a diversity which span a spectrum of activities from constructive play to industrial applications. The manpower to fuel this strategy, to maintain our lead in related areas of science and technology, and to begin to address some of our deficiencies will be integral to its success. With the changing demographics of our 18 to 24 year olds (decreasing by 30 percent by the mid 90s) and with the increasing proportion of minorities in the pool of students, the National Research Network will become even more important, allowing potentially all American universities to participate in or to take advantage of high performance computing research. Over time, our high quality university infrastructure will be broadened and extended as critical tools for research no longer have to be co-located at an institution.

NETWORKING

In the early days of the telephone system, several companies competed to deliver service. I am told that households often had several phone sets, one for each company which served them. Each set had a different number and the competing networks did not interconnect. Common directories did not exist and some people could not communicate with others. The situation today for U.S. research scientists using computer networks is somewhat similar to those early days of the telephone system. Agencies have created their own network (or in some cases several networks) to support their individual missions. The interconnectibility among these networks is poor, protocols are often different, and a network user may have five or six different identification numbers. A typical science building on a campus may have four or five connections to different networks and the various users may not know about one another in the same university department. There is usually no common directory and the best way to obtain an identifier is to ask an individual user. The maximum amount of information sent between users is usually severely limited, promoting mostly trivial conversation or at best electronic mail. Real time connections to advanced computers can be frustrating due to wait times or unexpected disconnects. This is discussed in detail in the National Academy report. One of the ironies of producing the network report to the Congress for the FCCSET Committee members was that the most reliable, fastest way to exchange drafts was via Federal Express (and we had the highest priorities possible).

The OSTP report recognizes that today's science is often not contained in single isolated disciplines. Breakthroughs frequently depend on cross-disciplinary research and greater degrees of collaboration, often occurring across research institutions and with closer involvement of our high technology industry. The National Research Network will increase the accessibility of geographically dispersed universities to supercomputers, allowing significantly greater numbers of researchers to address science and engineering problems of value to our nation and to benefit from the fruits of high performance computing research.

Fortunately, a National Research Network can build upon the nearly two dozen existing low-speed agency networks and make use of several national commercial fiber optic networks. A subset of the FCCSET committee called the Federal Research Internet Coordinating Committee (FRICC) has initiated a process of combining agency networks through memoranda of understanding between the National Science Foundation's NSFNET backbone which connects their supercomputer centers and DARPA's ARPANET, DOE's MFENET and NASA's NASNET. With the NSF acting as the lead agency in the internet, a higher speed 1.5 Mb/sec backbone recently began operation. Many research universities in the nation can now link to this backbone via regional networks typically operating at 56 kb/sec but individual users are normally locally limited to 1.2 or 9.6 kb/sec. We see this three tiered hierarchy (a national backbone, regional networks, and local institutional networks) evolving over the next few years to reach out to most educational institutions and moving to higher capacity in support of advances in collaborative research. It has been proposed that this association become the National Research Network and that it undergo three transformations as described in the November 20, 1987 report:

PHASE I would consist of several 1.5 Mb/sec backbone lines, regional networks which operated at 1.5 Mb/sec, and institutional access at 56 kb/sec (with selected institutions having access at 1.5 Mb/sec). The task ahead includes additional interagency collaboration, the harmonizing of protocols and the sharing of transmission facilities. I would expect this phase to be nominally complete within two years.

PHASE II will be an upgrade of the internet to provide 1.5 Mb/sec access to 200 or 300 research facilities across the nation; some of which are anticipated to be industrial laboratories. This phase will require several 45 Mb/sec backbone lines. Development and installation costs will be $5 M per year and operating costs will be perhaps $40 M to $50 M per year. This phase should be nominally complete within 5 years.

PHASE III is much more ambitious and requires substantial R&D before it's implementation. This phase would provide 1.5 Mb/sec service to nearly every research facility, 45 Mb/sec to approximately 1000 research sites, and 1 to 3 Gb/sec to selected sites with the backbone correspondingly upgraded. It is estimated that Federal development costs will be $400 M spread over ten years for this phase and operating costs will be $200 M/year. Optical switching, traffic control, advanced security, and new protocols will be necessary before it can be implemented. The committee on Computer Research and Applications estimates this phase to be deployable within 10-15 years.

Obviously, any incremental funding for networking must be evaluated and approved within the context of current activities and research needs in other high priority fields.

OPPORTUNITIES

When this study began, the FCCSET committee and members of the congressional staffs thought the overriding reason for a National Research Network was to provide access to supercomputing centers for the bulk of American university faculty and students who could not participate in research at the frontiers of high performance computing. It was learned that an equally important reason is collaboration; the sharing of software and databases, and the natural association of geographically distributed but functionally complementary individuals. The committee on Computer Research and Applications believes that one day researchers will as easily share research with colleagues via the National Research Network as they talk today with the telephone. A third important justification for the National Research Network lies in the innovation and productivity provided by a common quality network. We have researchers today working on distributed hypercubes (where individual computers act as modes with a massively parallel array of other--perhaps distant--computers). In a blue sky vision, I see the prospect for one day wheeling computing power around the country in much the same way that we wheel electrical power over the power transmission grid. And with a semi-intelligent network, I see the prospect of launching "knowbots" onto the national system of databases which collect information and return home with a memory full of information much like a robot would return with pieces for assembly. Of course research on the management and control of the system to insure the benefit and minimize misuse is essential.

BUDGET

Although the returns to U.S. competitiveness provided by this strategy will require additional investments by industry, states and the Federal government, the committee on Computer Research and Applications has made sure this is not just a more-of-everything approach. The economies of scale provided by combining several agency networks have already allowed for the implementation of Phase I of the National Research Network within existing resources. The long-range payoff to the American taxpayer will not merely be an improved balance of trade in the sale of computers because of the leveraging that high performance computing produces in other technologies of economic significance to the nation.

However, to provide these services within the proposed timelines, concept definition and R&D must begin now. The FCCSET committee is beginning the coordinating of high performance computing budgets in consultation with the Office of Management and Budget through a cross-cut budget analysis for the FY 90 budget cycle. This analysis will provide a realistic coordinated multiagency approach within the priority setting process of individual agencies. Dr. William R. Graham, the Director of the Office of Science and Technology Policy, has just sent a memorandum to all agency heads confirming his support for this coordinated approach.

I am aware that this approach to a computing strategy will be difficult to achieve because it requires so many ongoing partnerships, but it also has many strengths. It serves the U.S. in a geographically distributed way, it does not require a big science vs. little science decision, budgets can be stable, and agency priorities can be met. We believe that with the support of congress, the enhanced computing and information environment of our scientific community will produce enormous paybacks.

Senator GORE. Thank you very much. I am going to hold off on questions until Dr. Kleinrock has completed his statement. I want to call on you now and welcome you. Please proceed.

STATEMENT OF DR. LEONARD KLEINROCK, CHAIRMAN, NATIONAL RESEARCH NETWORK REVIEW COMMITTEE, NATIONAL RESEARCH COUNCIL, AND PROFESSOR OF COMPUTER SCIENCE, UNIVERSITY OF CALIFORNIA

Dr. KLEINROCK. Thank you very much. I am very pleased to be here, Mr. Chairman and distinguished members of the Science, Technology and Space Committee. I would like to address some of the nation's networking needs and mainly to summarize a report that the National Research Council put together reviewing one portion of the OSTP report, namely, that on the National Research Network proposal.

Let me begin with a few comments. I think a critical threshold has been crossed which has largely gone unnoticed. Two powerful technologies, data communications on the one hand, and data processing on the other, have finally converged in a way such that they will never again be separated. You cannot talk about one without the other.

Therefore, it is important that we think about that interaction. One of the results has been that our scientists have available on their desktops today enormously powerful workstations. And in most of their facilities they have very large centralized machines, databases, and software packages to which they need access.

To connect these two (the scientist with his workstation and his remote machines), we provide networks. Unfortunately today's networks are seriously inadequate to match the power of those two kinds of systems which we want to connect.

Right now networks lack sufficient and significant connectivity. Reliability is important. The heterogeneity problem is severe. The responsiveness of the network is not adequate. User-friendly access and dialogs do not exist. User support is lacking across many different kinds of systems.

The result is that the researchers' productivity is suffering and suffering badly. We were able to get firsthand evidence of the researchers' frustration through the surveys listed in our Committee report. The sad thing is, ours is the country that pioneered networking 20 years ago and we are now beginning to slip in our leadership and in our advanced networking technology.

Our Committee felt that the creation of a National Research Network is a superb idea and we strongly endorse it. We found enthusiastic support across the entire research community: physics, mathematics, chemistry, astronomy, astrophysics, oceanography, all fields of engineering, life sciences and medicine. They all have very enthusiastic and important needs for advanced networking.

The National Research Network could significantly enhance the productivity and quality of output of U.S. research across these many disciplines. As an example, in a recent review of the five national supercomputer centers, it was recognized that the major technical issue confronting them was networking and telecommunications. And it was unanimous across the five centers.

Improved science and technology is needed to improve competitiveness and productivity across these disciplines in the U.S. economy.

OSTP suggested a three-phase structure and we also support that structure, the first phase being interconnection of existing networks, the second extending support to far more institutions than are currently supported and the third phase, the more interesting one in our opinion is a dramatic increase in speed and function using advanced network technology.

The first two phases simply need coordination and good management. It is the third phase that needs a strong research program, not only in terms of providing higher speed, with which we are all familiar, but in terms of network management and control, in switching, routing, multiplexing, processor interfacing, protocols, connection-oriented procedures, and all those many technical issues that, indeed, become apparent and severe when we go to very high-speed networks.

We need basically a clean sheet approach to understanding how to deal with very high-speed networks and these issues have not yet been dealt with even in the research community in the depth that is needed.

Senator GORE. In this country?

Dr. KLEINROCK. In this country.

Senator GORE. Japan is dealing with them right now?

Dr. KLEINROCK. They are dealing with the high speed but they are not addressing some of these higher level issues like the protocol issues and heterogeneity issues. The Japanese certainly present a serious competitive threat here. They are putting more money into fiber optics than they ever did into semiconductors. They believe, that in the future, computing will be done with silicon and communications with fiber.

Senator GORE. Well, several years ago they identified as a national priority project the development of a ten-gigabit, ten-gigabit network.

Dr. KLEINROCK. And they are moving very rapidly in that direction. They are definitely taking our technology and turning it back on us.

In order for this National Research Network to succeed, we need a stable source of funding, we need a proper management structure and we need to put in place the research effort now to get that Phase III in place and in time.

There is no one agency that is a natural lead agency for this effort since this is such a highly interdisciplinary activity. We need to introduce user charges which represent a change in the way research is funded.

We must introduce economic incentives to encourage industry to provide the existing technology for national research. Long-range planning to allow the transition of the NRN technology to industry is needed as well.

The commercial world has decades of experience in network research and in providing accessible and reliable service to millions of users. Whatever form the National Research Network takes, it must recognize it has to accept and fold in the technology of the existing large scale networks that exist, including the commercial carriers.

Effective management is necessary. We need an experienced research executive and oversight committee, and we need funding through normal government channels. What came through very clearly in our deliberations was that the user groups feel neglected and unsupported throughout the entire network community right now. They feel that their needs are not being met, that they are not being heard, and that they are not being informed. Networks are being pulled out from under them. Their needs are being ignored.

They feel the technocrats have taken over networking as opposed to providing a service to the users and in our report we make it very clear that that has to be a prime consideration in any development of a network.

Our Committee endorses the OSTP's National Research Network as a general direction. Our concerns are the ones I mentioned, namely, technical considerations, funding, management and the folding in of commercial and specialty networks as they exist now.

The summary statement is that the challenge is to manage progressive introduction of increasingly sophisticated technology without compromising the user's ability to conduct his research for which the network is a tool. One must recognize that it is a tool. To carry that off is going to require substantial investment and that whole investment issue has to be considered very clearly.

We have a unique opportunity to boost U.S. competitiveness by providing this networking infrastructure. If the NRN is successful, it will ensure our leadership in science, research and technology and hopefully will allow us to enjoy the fruits of that leadership. Thank you.

[The statement and questions and answers follow:]

STATEMENT OF DR. LEONARD KLEINROCK, PROFESSOR OF COMPUTER SCIENCE, UNIVERSITY OF CALIFORNIA AT LOS ANGELES AND CHAIRMAN, NATIONAL RESEARCH NETWORK REVIEW COMMITTEE, COMPUTER SCIENCE AND TECHNOLOGY BOARD, NATIONAL RESEARCH COUNCIL

Mr. Chairman and distinguished members of the Senate Subcommittee on Science, Technology and Space, I am pleased to testify before your subcommittee on the nation's networking needs and to summarize the recent report of the National Research Council's committee (the National Research Network Review Committee—NRNRC) which reviewed the networking recommendation put forward in the OSTP report, "A Research and Development Strategy for High Performance Computing." I am Professor of Computer Science at the University of California at Los Angeles, and am Chairman of the NRNRC.

A critical threshold has just been crossed without much notice: in the United States today, data communications and data processing have converged in a way that now makes them inseparable. Unfortunately, however, the enormously powerful desktop workstations and centralized supercomputers used by the research community are currently being undermined by data networks that are seriously inadequate for this community's data communication needs. Current networks are not only deficient in bandwidth, they are also lacking in connectivity, reliability, heterogeniety, responsiveness, user-friendly dialogues, user support, and so forth. The researcher's productivity is being severely hampered by his inability to effectively link his high-performance workstation to the supercomputers, to the large scientific databases, and to his colleagues in a fashion that allows him to incorporate into his research these other resources. What is needed is an accessible, user-friendly, high-speed National Research Network (NRN), such as that proposed in the OSTP report, to provide these connections. Currently, however, this need is not being met and what the science research community sees instead is a serious mismatch between today's high-performance processing engines and low-performance computer networks. This situation is extremely unfortunate for the United States, the nation that spearheaded networking technology 20 years ago!

The creation of a National Research Network (NRN) is a superb idea which rectifies this shameful situation. The NRNRC found enthusiastic support fot the concept of an NRN across a wide specturm of researchers. That committee strongly endorsed the vision and promise of an NRN, recognizing that it could vastly enhance the productivity and quality of output of the U.S. research community, already a world leader. In several fields (notably the physical sciences), the progress of research is crying out for state-of-the-art networking services; indeed, referring to a recent survey of the five National Supercomputer Centers (reference: Testimony of Dr. Clifford E. Rhoades, Jr. before the U.S. House of Representatives, Subcommittee on Science, Research and Technology, Committee on Science, Space and Technology on March 24, 1988), Dr. Rhoades point out that, "The major technical issue that appeared in several of the reviews concerned networking and communications." The need for networking reflects the proliferation of information-processing capability in many disciplines. This multidisciplinary environment raises a number of grand challenges whose solutions are attainable, in large part, through the networking capabilities recommended for the NRN; such problems were hitherto out of reach. While the OSTP proposals emphasize goals for high speed, the committee felt that of equal importance were the concerns for widespread access, user-friendly service orientation, quality, and performance. Enhanced networking is absolutely critical for maximizing the use and benefits of supercomputers, but it is also essential for other scientific applications, something that can easily be forgotten in the excitement about supercomputing. Current and planned upgrades to NSFNET will provide some immediate relief for users of the NSF-funded supercomputer centers, but NSFNET alone will definitely not satisfy the large number of researchers who need computer networking to serve a large variety of needs.

The benefits of an NRN extend well beyond the research community. First, research is now, and always will be, a crucial part of the nation's business. By furthering the conduct of research and the transfer of research to those who develop and use technology, an NRN could greatly enhance U.S. economic competitiveness in a multiplicity of disciplines. In addition, the specific-results of the research required to achieve the NRN will advance the data processing and data communications fields significantly. That research will likely produce commercial spin-offs whose impact will extend well beyond the research community. Historically, networks developed for researchers (most notably Arpanet) have provided the foundation for commercial data communications networks. Without an NRN, the market might eventually get around to providing the same capabilities, but not soon enough to meet critical national research needs required to keep the United States competitive.

The committee supported the three-phase structure proposed by OSTP. In the near term (late 1980s), the NRN would interconnect the existing networks that are fragmentary, overloaded, and poorly functioning. In the second phase (in the early 1990s), it would extend computer network support to researchers at far more institutions than are currently served (in particular the smaller institutions not now served well) and it would significantly enhance the range and quality of network-based services provided to researchers. In the final phase (late 1990s) advanced technology and techniques would be deployed to provide a dramatic increase in speed and function to match the power of computers in use at that time. The technical demands for the first two phases identified by OSTP are not leading edge items. However, achieving the high-performance network carrying the high volumes of traffic envisioned in phase 3 for the late 1990s requires computer and communications research using a clean-sheet approach that must begin now. Among these critical research topics we include areas such as network management and control and appropriate technology for switching, routing, multiplexing, processor interfaces, protocols, connection-oriented communications, and layered architectures, all of which offer opportunities for technolgy transfer to industry from the necessary research program.

However promising are the prospects of an NRN, our committee recognized and emphasized that thorough planning and proper funding are necessary to convert this concept into a reality and a success. Funds for advanced development should be kept separate from funds allocated to running the operational network. The funding, managing and implementing of a sophisticated, capital-intensive infrrastructure presents a very different paradigm than that of traditional science research projects. One reason is that an NRN cuts across and supports all of the traditional disciplines in science. Consequently, no one agency is natural as the lead agency for an effort of this sort. Another reason is that to avoid the problems that have plagued smaller, less-sophisticated networks such as Arpanet, a transition to some sort of user charges appears to be necessary. We are talking about a far-reaching change in the system of research funding in this country. That change will take place in an environment of concern among researchers about the adequacy of funding for their research, per se.

To encourage the development and use of a widely accessible network service, the committee suggested that the NRN might provide (1) a universal basic service at some low to moderate speed for electronic mail at a low cost to users and (2) higher levels of service with an appropriate charging structure. A low charge for basic service would lessen the initial resistance that might otherwise be felt by first-time users and yet discourage the overloading that occurs with free networks. A charging structure would also make possible an eventual transition of the service to the private sector. For this to work, the research grant process must be changed to accommodate network charges, and other provisions would be necessary to sustain users from poorly-funded disciplines. New mechanisms, such as vouchers for network services, should be considered.

The committee recommended that economic incentives be provided to encourage industry to provide existing technology to the NRN, but that the government provide funding for advanced network development in cooperation with industrial research efforts where possible. It also encouraged long-range planning that would allow for the eventual transition of at least portions of a government-developed NRN to industry, because meeting the service and access goals of the network implies resources beyond the apparent capabilities of federal mission agencies or the National Science Foundation. The commercial network suppliers have decades of experience in network research and in providing accessible and reliable service to millions of users; it would be prudent to leverage this experience for the NRN. One might consider an NRN would be a stand-alone network, but the committee concluded that the proposed networks must build on existing and anticipated commercial network facilities and network-based services.

Finally, good, professional management is absolutely essential for maximizing the effectiveness of the investment in an NRN. Remember that we are talking about a system that could extend access to what could eventually be almost a half-million users in the research community. Management requires early (and continuing) attention and action, and it is a nontrivial element of the total cost of providing network service. The overall problems are to engineer facilities and services to match demand and tune them to the expressed needs of users, to maintain them in an operational state, and to operate them economically. To satisfy the funding and management challenges presented by an NRN there should be a funding mechanism through normal government channels, an oversight committee, and an experienced research executive from the private sector for day-to-day management; in addition, one or more user groups should be established to provide imput from and output to the user community.

In the remainder of this testimony, I elaborate upon the charge of the NRNRC, it procedures, and its concerns, conclusions, and recommendations. The National Research Network Review Committee (NRNRC) was established under the auspices of the National Research Council's Computer Science and Technology Board (CSTB) to review proposals developed by the Committee on Computer Research and Applications of the Federal Coordinating Council for Science, Engineering, and Technology (FCCSET) for the establishment of a high-performance, national computer network for researchers. Through CSTB, the committee was asked to consider three sets of issues:

The technical feasibility of the network proposals.

The utility of the proposed network to the research community.

Devlopments in computer technology that might impinge on the proposed network and the services that might be offered over it.

Our report primarily addresses the network proposal presented in the final OSTP report, but it also draws upon background material conveyed in earlier, more detailed drafts from FCCSET (Committee on Computer Research and Applications, FCCSET 1987, Reports from the Workshop on Computer Networks, February 17–19, 1987, San Diego, California, and A Compendium of Supporting Technical Data). One result of the changes leading to the OSTP report was a shift in the committee's emphasis away from specifics of the network proposals and toward a more general discussion of the concept of and parameters for a high-performance computer network. Because of its concern with the value of the proposed network to the user community, the committee also sought imput from representatives of several scientific disciplines both through direct contacts with researchers and through the use of electronic bulletin boards.

The committee viewed the OSTP proposal as providing a general direction for the NRN, not a set of design specifications (which it was not) to be evaluated. Four areas of concern were addressed: technical considerations, funding, management, and commercial and specialty networks, which have been highlighted in my remarks. The challenge underscored by the committee is to manage the progressive introduction of increasingly sophisticated technology without compromising the user's ability to conduct the search for which the network is a tool. This may not be possible without substantial investment, and, even with it, compromises may be necessary.

We have a special opportunity with the NRN. We can give an enormous boost to U.S. competitiveness by providing this networking infrastructure. But we must not hesitate, for our foreign competitors are busy at work with their own national research networks (e.g., the RACE program in Europe, and national network projects in Canada, Japan and in newly industrializing countries such as Singapore, and so on). If we are successful in the NRN development, our competitive advantage will continue to grow as this unique infrastructure dramatically improves the quality of output and productivity of the research community of this nation.

QUESTIONS OF THE CHAIRMAN AND THE ANSWERS THERETO

Question 1. Computer viruses and the dangers they pose have been getting increasingly more publicity. Because computer viruses can spread so easily through a computer network, the National Research Network could end up increasing the threat posed by computer viruses.

How much of a problem do viruses pose to researchers on the National Research Network?

Did your committee consider the problem of computer viruses?

Do you think that there will be many researchers or businesses that will be reluctant to hook up to a National Research Network because of their concern about computer viruses?

Answer 1. Our committee certainly feels that computer viruses do pose a threat to research which would take place on a national research network. There are, however, certain protective procedures that a conscientious researcher can take in order to protect his information, (e.g., saving various back-up versions of his material as he progresses through his research). This will protect him from losing all of his data but certainly will make it possible for the recent portions of his data to be corrupted if a virus were to attack his system maliciously. The more insidious danger, however, is in the case in which a virus changes programs and data in an undetected fashion. This, indeed, could be worrisome to certain types of research taking place. However, it is important to note that one of the major applications of the National Research Network would be to get to archival and encyclopedic data which should be protected from any changes by means of an access code; this data would then be protected to a high degree and still be of extreme value to a national network of researhers. In order to bring about the protection needed for that type of application, we must begin the planning now.

Our committee, in fact, did consider the importance of both security and privacy issues and concluded that these issues are currently underemphasized, leaving our networks underprotected. We specifically did indicate (pg. p. 33 of our report) the dangers of viruses and other security risks. These dangers represent only one aspect of the difficulty of implementing large multi-faceted networks which are responsive to users.

I believe a small fraction of the research community would, indeed, be reluctant to hookup to a national research network because of their concern with computer viruses. Certain business organizations will also feel reluctant due to the proprietary nature and mission-critical nature of some of their data; however, they certainly can protect themselves easily by the use of stand-alone and/or special-purpose terminals.

Question 2. There is an inherent tension between computer scientists who want to have a high-speed network to experiment with and computational scientist (physicists, engineers, and others) who need networks to access supercomputers. Computer scientists who study networks want state-of-the-art technology. Other scientists may not need a state-of-the-art network, they just want it to be reliable.

Is it possible to resolve this conflict?

You had representatives of both camps on your committee, could they come to agreement on what a national network should look like?

Answer 2. Regarding the conflict between computer scientists and computational scientists, I do not believe that a severe conflict exists. Our approach was to recommend a reliable and modern network technology for the use of computational scientists which would be stable. This network would not be the laboratory for experimentation by the computer scientists. Thus, there would be no conflict of the sort with which you are concerned. As regards the form of the national network in its advanced stages (phase 3), our committee felt that a strong research program must be initiated now in order to decide on the form of that eventual network since, at this point, it is unclear what the critical technological choices should be in order to bring about an advanced National Research Network.

Question 3. I feel strongly that a national network could help spawn new businesses and speed up the industrial application of new technology.

Did your committee discuss how much access researchers from industry should have to the national research network?

Do you think that industrial users should help foot the bill for the network?

Answer 3. The issue of extending the national research network to business and industrial applications is an interesting one. I believe that the extension in this direction is a very significant one and is important but will require a much more extensive and carefully thought out effort than the one planned for the National Research Network. The concept of such a National Information Network has already been addressed in a committee of which I am a member, namely, the Computer Science and Technology Board of the National Research Council. That board has already formulated a plan, which has recently been approved, to study, among other things, problems in the security area for such a National Information Network. It is a natural but fairly elaborate step to go from the Research Network to the Information Network. If, indeed, industrial users were to use the National Research Network then certainly they should help pay for the cost of that network.

Question 4. Your report encourages those in charge of a national research network to develop long-term plans for phasing out the government-run network and letting the telecommunications industry meet researchers' networking needs, just as they provide telephone service to researchers today.

Was there general agreement among your committee members on this point?

How long do you think this transition will take?

Answer 4. Regarding the phase out of the government-run network and letting the telecommunication industry meet the researcher's networking needs, there was fairly general agreement among our committee that this would be a good move. It is likely this transition will require roughly five years to come about, but the important thing to note is that we bring the telecommunications industry in at the earliest possible point to discuss the technology, phasing, management, and financing of such a network.

Senator GORE. Thank you very much. I have a number of questions and even though we started a little bit late, I want to take some time today to do a thorough job on this subject. I felt for many years that this was one of the real keys in positioning our nation for the information age.

Dr. Huray, two years ago I referred to the Congress legislation that required OSTP to review the networking needs for supercomputers. The report that OSTP released last November considers the nation's networking needs but also contains a broad and ambitious proposal to double the amount of funding for the entire field of high performance computing.

Why did OSTP decide to expand the scope of its report? I am glad it did but can you tell us how that came about?

Dr. HURAY. Yes, sir; I can. I might note also that we recommended in our budget to triple the amount of funding in networking by itself over the next five years. We do think, however, it is important to address this problem in an integrated fashion. We think that a network, a National Research Network, should transmit good science, not just electronic mail, but provide software exchange between participants. We think hardware should also communicate well with the network.

We wanted to take an integrative approach, therefore, and so we restructured the Committee to carry out the report. One of the net subcommittees dealt with just networking and it focused upon the report to the Congress, as you requested. We produced three volumes for that report (which were not entirely internally consistent, I might add, and, therefore, hold no official status). We also produced four other volumes, however, by two other subcommittees, one which dealt with research and development and one which dealt with applications. We felt it was important to put the National Research Network in the proper context to make sure we had perspective among other computing opportunities. And we wanted to set some priorities. That is why we came forward with a balanced approach.

Senator GORE. Now, networking, research and development and applications.

Dr. HURAY. And making sure we have the manpower (human resources) of course.

Senator GORE. In order to cover the waterfront, would you not have needed one in artificial intelligence?

Dr. HURAY. We chose not to study several areas, including artificial intelligence. We also did not study office automation or the manufacturing of compoments and left out a number of broader issues. We wanted to stick fairly close to high performance computing because we felt that it will be the activity for which the network provides the most service.

Senator GORE. Well, in microcosm I think we can see the kind of synergy that comes from networking in the dialogue between the Legislative and Executive branches on this question.

We initiated an inquiry into networking. You took that initiative and properly expanded it to look at a more comprehensive approach and as we receive your reports, we are going to expand our scope and give the same level of attention to the other areas that you have explored. Then we are going to add a couple of our own and throw the ball back into your court once again. That is the kind of communication that takes place between researchers who have the capacity to interact with high data flows on a regular basis.

The $1.75 billion proposal contained in the report is quite surprising coming from an administration that has been famously reluctant to start big new programs.

Can I assume that the Office of Management and Budget has signed off on this recommendation?

Dr. HURAY. Yes, sir. This is a report of the Executive Office of the President and therefore OMB does back this report. Of course, I do not speak for the Executive Office of the President, in my appearance here today, merely for the FCCSET Committee on Computer Research and Applications; the group that did the study.

However, in our budget cross-cut analysis we will work with individual agencies to make sure that the priorities are set internal to the agency as well as in a larger umbrella program with the Executive Office.

Senator GORE. The payback on investment will be larger from this expenditure than any other investment that the Federal Government is likely to make.

If it is so important, why did OSTP wait until the final year of this administration to submit this proposal? I mean, I suspect it just has taken that amount of time to put it in perspective, to get around to it.

Dr. HURAY. Senator Gore, we wanted to do a complete study. As you recall, the final legislation, Public Law 99-383, called for two phases of the study, the first dealing with supercomputers, due within one year of August 20, 1986, and the second dealing with all research computers in a second year; that would have been due this year. We decided to not address those issues separately because we felt we could not differentiate the network needs for supercomputers from other research computers. We did submit the report earlier than the second year (which would have been due August 20, 1988), but slightly after our first year deadline of August 20, 1987.

Senator GORE. The Congress is frequently lobbied by groups of scientists who would like to see increased funding for their particular fields. Most of the participants on the FCCSET Committee you chaired are computer scientists who rely on supercomputers. Cynics might say no one should be surprised that a panel of computer scientists would propose spending more money on computer research.

Dr. HURAY. Senator Gore, this study was carried out with the participation of hundreds of researchers across the country. Although the participants that are listed on the last page of the report are the official representatives on the Committee, there were five separate hearings held across the country with participation of not just computer scientists but of other kinds of research scientists. They involve industry and universities as well as Federal Government. However, this strategy is a multi-agency Federal Government strategy and we do still have to work with industry and universities to refine it.

Senator GORE. The entire scientific community benefits from this, not just the computer science community?

Dr. HURAY. Yes, sir.

Senator GORE. We have heard a lot from Frank Press and others about setting priorities in the sciences. We do have a limited budget. We also have an enormous deficit that has to be brought under control.

Which aspect of computer research most needs additional funding right now?

Dr. HURAY. Well, we put them into priority standing by giving you balanced recommendations with recommended funding and, therefore, they are budget prioritized. We would not like to say that one aspect should take place at the expense of the others.

Senator GORE. If we adequately fund the network initiative and create the synergy, for lack of a better word, and access to databases, increased research and development capability, how can we and should we prevent the benefits of that increased capacity from being shared with every researcher, commercial or educational, in the entire world who can manage to hook into the network?

Dr. HURAY. As you pointed out, policy lags even software. We have not made a policy decision about how we will approach that problem. As Dr. Kleinrock has pointed out, there are still many areas of research, protocols, security and policy yet to be done.

The National Bureau of Standards, for example, is participating in this to help us with some of the technical aspects such as protocols and security, routing, traffic protection, other kinds of protection of databases, but we think some policy decisions have yet to be made relative to the international community.

Physically, international access is easy to accomplish. In fact, some of the existing networks already interconnect internationally. So we think with the help of Congress we should address these problems but we have not yet done so.

Senator GORE. Is it physically easy to prevent East German kids from playing in Los Alamos laboratory's database?

Dr. HURAY. Well, sir, as you know, that has already happened and with other university students connecting, we expect interesting things in the future.

Senator GORE. The Federal Government runs dozens of different computer networks. NSF, NASA, DOE, DOD and others all have important national research networks of their own.

Some are interlinked, others cannot be linked because, as you all have indicated, they use different protocols to encode and transmit data and there are all kinds of problems.

Are there efforts under way to standardize Federal networks so they can be linked other than the recommendations of this report?

Dr. HURAY. There are activities occurring which memoranda of understanding are being written between agencies to interoperate our networks and those are resulting in protocol changes. Some specialty networks probably will not participate with what we eventually call the National Research Network.

In fact, I am not sure we will ever say we have a single network. We think there always will be specialty networks. There will certainly be networks which operate for the purpose of research and development, looking at the very high-speed capabilities, operated by scientists who regularly bring the network down to test its capabilities. And so we see the need for several networks in the future. I am not sure if I answered your question.

Senator GORE. You did. There is often a tendency to overemphasize the very highest high technology, the most powerful computers, the one-of-a-kind devices that capture headlines and yet often the biggest impact comes from lower technology devices that are manufactured by the millions.

Somebody has said Americans get the Nobel Prize and Japanese get the profits.

Is it possible that the OSTP proposal overemphasizes the big and the flashy at the expense of the ordinary and everyday and most marketable?

Dr. HURAY. I hope not. We, of course, have emphasized the opportunities, I tend to do that personally, because it is an exciting area to be in. But we are proceeding in a very prudent fashion to bring together existing networks to produce a very reliable system which is easy to use by our researchers across the nation. We hope that we can address individual complaints of researchers as they begin to feed back their needs on the production network.

Senator GORE. Supercomputers will come down in price fairly rapidly; is that safe to predict? I mean they already have but that is likely to continue and accelerate, would you say that?

Dr. HURAY. Yes; prices of supercomputers will decrease but capabilities will proportionally increase. I am not sure which will happen. Probably both.

Senator GORE. Probably both? What is the average cost of a supercomputer today?

Dr. HURAY. Well, of course, we could give you the price of some particular manufacturer's supercomputer. Let's say they are certainly in the range of millions of dollars for the highest performance, single-clock computers.

Some of the more experimental computers, which operate with massive parallelism, are substantially less than that.

In fact, we find those often on university campuses and in Federal laboratories. We would like to see additional distribution of those experimental computers for people to effectively play with new architectures, develop innovative new schemes, and feed ideas back into the industry. It is safe to say that some of the costs for experimental systems are quite within reach for even university departments.

Senator GORE. As those costs drop, the number of machines in use will expand, will increase dramatically, the need for a network of the kind we are discussing today will increase dramatically.

Many people in the telecommunications industry feel that industry can meet the networking needs of researchers. AT&T is selling tens of millions of dollars worth of networking services at 1.5 and even 45-megabit rates and it has even installed a small gigabit network for one agency that shall remain nameless.

It would seem that the Phase II network that your report proposes to install by 1990 is not that far removed from what they claim is available in the commercial industry. What is the response to that?

Dr. HURAY. Well, as you know, and as Senator Riegle correctly pointed out, IBM and MCI are currently participating in the existing Phase I backbone for the NSF centers. We expect more participation, including from AT&T and other network delivery commercial systems in the second and third phases. In fact, we want to make sure that it happens. We think it is important to have that kind of partnership.

However, it is important to recognize that commercial network systems rely upon profits, and profits in the area of research are probably not comparable to those in ordinary telecommunications, such as in the area of telephone service, so we think that those commercial carriers may place other things at higher priority than a research network. We think it is important, therefore, to form this partnership and have Federal and industrial investment in a combined network.

Senator GORE. So you have commercial incentives to proliferate networks with relatively low data rates and at the same time research needs for a much larger capacity network.

Putting in place the larger capacity network, one creates some anxiety on the part of the commercial participants who stand to gain continued profits from the small data networks.

How do you allay those fears?

Dr. HURAY. Well, we think it will be to the long-term advantage of those commercial carriers to work with the Federal Government in bringing up the networks because it will, in a sense, create a new market. As I said, I think one day the culture will change so that scientists will easily communicate with one another over a high-speed network exchanging data, and that will lead to a new kind of information transfer which, of course, will have profits associated with it.

In the early stages, those profits may be small. We think in the long run it will be to the benefit of those industries to participate with us.

Senator GORE. Dr. Kleinrock, you alluded to this new age where research made possible with the high-capacity network will yield increased profits and productivity. What did you have in mind there?

Dr. KLEINROCK. Well, certainly this country has a long history of reaping profit and benefit from the research efforts across the scientific community. We see every reason to believe it is going to continue.

The spawning ground for much of the science research in the universities and in the research laboratories across the country are these groups that are asking for the connectivity which will hopefully provide the interaction, not only, as you said, among the databases but also among the researchers themselves to enhance and accelerate scientific progress across the board: in new materials research, in new graphics devices, in new technologies for oceanographic exploration, for oil exploration, et cetera. So there is clear commercial benefit from the science research that is taking place. We have seen that in the past and we anticipate it in the future as well.

Senator GORE. How do you measure it in its incipient stage?

Dr. KLEINROCK. That is difficult to do. Some of it is prediction, some is serendipity.

Senator GORE. It is impossible to do?

Dr. KLEINROCK. There is no reliable measure you can apply to predict profits from investment in infrastructure.

Senator GORE. That is the problem. I am fond of drawing an analogy between the interstate highway system and what we need to create with the supercomputer network.

In the case of the interstate highway system, the nation was able to telescope the investment of several decades into a few years because the users of the completed system could reasonably be expected to pay user fees over a long period of time that would make the project self-financing.

And the superhighway system as a result did not add a penny to the budget deficit, ever. The productivity gains which came as a result of the completed system have never been measured. They have just been enormous.

The same thing would occur with a supercomputer network.

How can we complete the analogy and find a way to collect user fees in some form from those who in the future will clearly benefit from this network? We have shortened the distance between the laboratory and the supermarket. And we can anticipate profits from research, but we cannot measure them and we cannot, as yet, figure out how to anticipate a user fee from the people who are going to get the most good from this network.

Is my question clear?

Dr. KLEINROCK. Yes, it is. In partial response to that, we recognize this network is an infrastructure just as is the superhighway system, and as is the telephone system. And by itself it is very difficult to measure the direct profit which may be derived from all of the ancillary by-products that come from it.

We have, in our review, suggested the introduction of user fees on the science and research community but, as you point out, the profits we are talking about are beyond that stage. These profits will eventually come from the industrial applications. It is not clear how we tax those people for the benefit they are going to derive inasmuch as we do not tax them now for the cost of educating the people who then finally work in their institutions and provide the products from which they derive profit.

Senator GORE. Now, the user fees you are talking about are not revenue-raising devices, they are merely intended to allocate capacity between different federally funded agencies and institutions, correct?

Dr. KLEINROCK. Yes, but beyond that it is meant as an incentive to draw in the commercial environment to provide these services and to replace the government funding and government support of the function.

Senator GORE. All right. So it would create a financial structure that could be a gateway for commercial users to share the network with Federal agencies and so it could indirectly have that purpose.

If it is impossible to collect the money through the kind of user fee in sufficient quantities to finance the project, and if we have the intellectual courage as a nation to predict the obvious benefits that are going to come from it, then the only alternative is to go to the general revenue and say: This is worth it, in spite of the fiscal environment, we are planting seeds, we are investing in the future, it has a high probability of enormous payback, so we ought to pay for it. That is really the answer, it seems to me, unless we can figure out a way to complete that analogy to the highway system.

Dr. KLEINROCK. I agree with the former statement you made. I think we do have to recognize it is an investment at this point in time.

We also must recognize that this network investment is in some sense small relative to enormous investment we have in the devices and researchers who will attach to and benefit from this network investment. This is a relatively small increment with the potential of a very high leverage.

Senator GORE. All right. Now, what about this interaction between the commercial networks that are in place and others that are in the planning stages that have a relatively low capacity and this proposed network? Let me give you a hypothetical example.

Let's say I am a company planning to invest in a five-megabit fiber optic cable linking Champaign-Urbana, Illinois with Oak Ridge, Tennessee and that I plan to sell capacity on that cable to commercial users and I have done a market study that tells me that I can recoup my investment in a short period of time and make a handsome profit.

I find out from the trade publications that the Federal Government is going to put a 3-gigabit line right next to the one I am planning to put into the ground and is going to experiment with user fees that allocate capacity to the research teams at Champaign-Urbana, Oak Ridge, Tennessee, but that also will entice commercial users to share that same capacity.

I am not going to like that. Tell me why I should feel comfortable with that.

Dr. HURAY. Let me say that I am fairly confident that it is no one's intention in the Federal agencies to lay fiber optic cable. We plan to, and have brought up, the existing national networks with the cooperation of private industry. We think that competition will be a very important aspect of that process, however.

As you know, when the bid was let for the merit program for the national backbone, there was a competitive process and it did result in industry coming forward with significant matching dollars and there were others there competing with them.

We think that competitive process should continue and that we will achieve profits for industry and maximum benefit for the Federal Government. The way we see it, it has to be a partnership.

Senator GORE. I understand that. But the hypothetical is intended to illuminate an underlying relationship that has to be worked out to the satisfaction of all involved.

Now, in effect, if the Federal Government is going to subsidize the extra 2 point something gigabits for the research needs, and then attempt to recoup part of the taxpayers' investment in those 2 point something gigabits, that effort to recoup will potentially come into conflict with the 5-megabit capacity that is purely privately funded, unless there is a clean separation and an accurate accounting.

Now, I am sure there is an easy solution to this but in communicating the need for this national network with private companies, I encounter a great deal of apprehension and I think it comes down to this relationship somehow. Am I wrong about that?

Dr. HURAY. No, sir. I think this still needs to be refined with the participation of industry and we have some work to do yet.

Senator GORE. What is the solution to it?

Dr. KLEINROCK. May I point out, we addressed this issue in part and we felt it is important to give the user, the science researchers here, real funds, perhaps special "network data rights" which equate to dollars, in which he can exercise freedom of choice as to where he buys his service. If the commercial offering is competitive with this government funding, go ahead and do it. We are not recommending that——

Senator GORE. But, let me interrupt briefly. You have still got this enormous disparity in the capacity typically required for commercial service and the enormously larger capacity anticipated for research needs.

Now, have I got that wrong?

Dr. KLEINROCK. No. But the research community is a customer with dollars to spend. It may be that the commercial sector will be slow to satisfy those needs, whereas the government could respond more quickly by taking the leading edge. The idea is that these networks will mature and eventually be taken over by the commercial sector. One way for promoting that is through competitive freedom of choice on the part of the user.

Senator GORE. It is only a matter of time before industry itself will be using these enormously high data rights?

Dr. KLEINROCK. That's correct. They will quickly react. We see a new mode of interaction that is heavily graphics and image orientated. We see that as an emerging mode of communication for the commercial environment as well.

Senator GORE. We have faced a similar challenge in the past and the COMSAT model resulted. Not everyone was happy with it and it was substantially modified before it was put into place. But there are models available that ought to be actively supported.

Has the industry's reaction been favorable to the OSTP report?

Dr. HURAY. There have been a number of reactions to it, many of them have been favorable. I think that some of the later panels are going to be addressing this issue, Senator. We have not decided exactly what management model will be used in operating a National Research Network; whether that would be a quasi-public corporation, a Government owned corporate operator, or a direct activity of a Federal agency, so industry has not had a uniform response. They have been making some recommendations to us and we are looking for the right forum to address this question jointly.

Senator GORE. I think we really need a lot more attention to that relationship. We need to explore potential models to make the relationship a healthy and creative one and to explore the potential for using the network to promote the transfer of technology out of labs and into the commercial arena.

Everybody realizes that one of our biggest challenges is in applying the results of our basic research. That is where the Japanese are beating us so badly. We do the research and then we watch them turn it into products. It may be that in finding a way to manage the network, we can find creative responses to that larger problem.

At the current time a majority, a vast majority of the Federal funding for advanced computing is going to military applications. Is that a healthy situation?

Dr. HURAY. Well, there are certainly applications of high performance computing in the military that is appropriate and satisfactory to their mission. Much of the research and development in the Federal Government is being carried out at DARPA. Most of it is not military-oriented but it is for the benefit of all science and technology. We think there is good science going on at DARPA and we support that activity.

Senator GORE. Does the technology developed there get put to adequate use in the civilian sector or does a lot of it just not see the light of day?

Dr. HURAY. I think there is a lot of commercial activity coming out of the DARPA programs. Several of our small industries have been spawned as a direct result of some of that innovative research.

Senator GORE. I am certainly aware of that. I just wonder if some of it does not. We will explore that more.

Moving to the area of international competition, the OSTP report warned that Europe and Japan are aggressively moving ahead of the United States in a variety of networking areas with the support of concentrated government and industry research and implementation programs.

I think we can readily imagine the economic consequences if we allow this to take place. There's a vivid analogy, and I do not know who said it first, between the steam engine in the industrial age and the supercomputer in the information age. The supercomputer is the basic generator of new wealth and if we fall behind in this technology, it will be to our peril.

Could you give us some examples of where the United States is falling behind, particularly in networking, but any other areas as well?

Dr. HURAY. In the early stages of our review we did examine several networks such as the Esprit program in Europe and a program in Japan involving industry and Federal laboratories. They were clearly moving ahead of the U.S. and, in fact, have in place several operating research networks. However, since the time of that report we have moved forward and we are beginning the implementation of our own network. We intend to lead eventually.

Dr. KLEINROCK. May I point out that the supercomputer is only one component of the information processing industry and other high performance, mid-range and low-end machines are also an important part of the fabric which makes up this industry. It is not only the supercomputers to which the researchers need access. There are workstations on the desk that are powerful. Powerful machines which are not necessarily of the supercomputer class provide enormous output in terms of research results.

Senator GORE. Well, the new breakthroughs in massive parallelism certainly modify the way we used to look at the role of supercomputers. Both the OSTP report and the Academy Study describe the problems with the present research networks and describe them as overloaded, hard to use, not as reliable as users would like, incapable of transferring large data files and strained with the task of handling the present volume of electronic mail. This indicates to me that networking is not receiving adequate funding. Obviously it is not. Why has it not up until now? Is it just the lack of any market signals or incentives or leadership or what?

Dr. KLEINROCK. I think in some ways it was a lack of recognition early on that these networks were to grow as large as they would. As a result, neither proper management nor adequate funding were put in place early enough. That, along with the heterogeneity of a collection of networks which initially did not know how to talk to each other, and in many ways still do not, caused the problem.

Dr. HURAY. Part of the problem is that each of those agency networks were created to solve a particular requirement. They grew and grew and then finally the need came to interconnect. There was no coordinating mechanism throughout those Federal agencies. We believe that the FCCSET process has played a very important role in providing some coordination and we hope to continue that coordination.

Senator GORE. Before the FCCSET initiative and the initiative of this Subcommittee, there was a lack of attention to networking.

The same lack of attention that led to these problems in networking is now clearly evident in the area of software. And while the market draws out software for some applications, is not producing other critical software as you said, for a variety of research uses that are not compensated by the market.

How do we fix that problem? First, do you agree that it is a major problem?

Dr. HURAY. I certainly do. As you know, that was one of the strong recommendations of our report; that we focus on particular activities such as generic tools and algorithmic research which would be useful to all of our researchers and shared among them.

We think it is important, however, to also focus on particular problems of interest to industry which will have direct economic payback. We think there ought to be Federal programs shared with our industry who are interested in certain applications. We would like to see two or three of those applications (or Grand Challenges) funded per year and we would like to begin to identify them jointly with industry. In fact, our FCCSET Committee has begun a process of trying to identify areas where we have real opportunity; where new computing power could be brought to bear on the solutions to particular problems and the results can occur right away.

Senator GORE. I introduced legislation several years ago to establish a national committee with the capability of providing a small amount of seed money for projects to produce software that was deemed critical for the country. If the application was approved by the panel, a research could then augment the seed money with a larger amount of private funding, attracted by the certification process. That process would identify software that is critical and thus likely to be marketable in the future. The administration strongly opposed that initiative.

If we had it a few years ago, I really think we would be down the road to solving this problem. I will provide the updated version of that bill and hope that the FCCSET Committee can convince the administration to turn around on it in the time after Labor Day and before the adjournment. Well, I said we were going to take our time and explore this carefully but time is running out and we do have two large panels to go, so we may well provide additional questions in writing.

I would like to thank both of you for the enormous amount of time you have put into this effort. I certainly do appreciate it, Dr. Kleinrock. And thank you, Dr. Huray. We certainly enjoyed working with you on this, and were very, very pleased with the response to the initiative that we launched asking for you all to look at it carefully. You both have done so and have produced excellent reports. I appreciate both the bodies of work. Thank you very much.

Our second panel, please come to the witness table, Dr. Robert Kahn, President of the Corporation for National Research Initiatives; Dr. Robert Dickinson, Deputy Head of the Interdisciplinary Climate Systems Section at the National Center for Atmospheric Research in Boulder; Paul Kutler, Acting Director of the Aerodynamics Division of the Office of Aeronautics and Space Technology at NASA; and Dr. Kenneth King, President of EDUCOM at Princeton, New Jersey.

Without objection, as I have already noted, your prepared statements will be included in the record. We invite you to summarize your statements if you can do so. We will begin with Dr. Kahn. Welcome, it is good to see you again and please proceed.

STATEMENT OF DR. ROBERT E. KAHN, PRESIDENT, CORPORATION FOR NATIONAL RESEARCH INITIATIVES

Dr. **Kahn**. Thank you, Senator Gore. I am pleased to have the opportunity to be here and to testify once again before one of your Subcommittees. The last time, as I recall, the topic was emergency management and I might point out that is a topic not entirely unrelated to the subject under discussion today. Increasingly we are seeing the importance of the research community in dealing with some of the short and the long-term interests of the nation that go beyond pure R&D. A National Research Network can help in that context too.

I plan to summarize some of the key points in my testimony and make a few brief observations. Many of the issues having to do with use of supercomputers will be addressed by others on this panel. I would like to focus on the concept of a National Research Net.

The benefits of networking to access supercomputers are fairly clear. People can visualize a wide range of possibilities right now, but I am willing to assert that we cannot yet even envision the most powerful applications of future very high speed networks.

I support very strongly the conclusion in the OSTP report that we ought to have a National Research Net. I think it is extremely important to broaden the participation in the use of research nets well beyond the computer science community. When we did the early work on the ARPANET it was really intended to benefit the computer science research community because there was no experience base at all in computer networks and not much knowledge even about time sharing computers during that time frame.

That era is well past, and I think it is now time to insure that all sectors of the research community can become involved with computation and the use of networks. In fact, there is a much broader issue lurking behind this whole discussion which has to do with the involvement of the rest of society in this whole process and that is an issue that you were, if I can say, alluding to, I am probably understating it. It is clearly a very important issue.

Senator **Gore**. Every science lobbying group that comes to the Capitol and lobbies for research money for this or that ought to have the research network on their agenda. I forget who it was who identified computer science as a whole new branch of science, a whole new way of investigating and learning about the world. This new way of thinking could affect the entire research community. So I strongly agree with you. We have got to get away from looking at it as a computer science project. Please forgive me for interrupting.

Dr. **Kahn**. We should consider you an honorary member of the technical community. You have been very helpful. The other point I would like to make is that besides broadening our view of what a research net is for and who should be a part of it, we need to develop our notion of what it takes to make such a network viable and useful. Simply supporting the ability to move bits around does not necessarily make a network useful.

The actual network users need support for the kinds of equipment (from workstations to supercomputers) that make it possible for them to use the network effectively, this will be especially important when we set about broadening our aspirations to other disciplines, whether they be physics or chemistry or biotechnology or whatever. Many of those research groups are not now outfitted in ways that will let them take advantage of networks from the point of view of equipment and software.

We are going to have to focus explicitly on helping those research communities learn to make effective use of networks, learn what the power of them may be. It is an incremental process that must take place over time. We are going to have to support them in that learning process as they begin to crawl and walk and then run. I think they can learn rather quickly.

There will, of course, be interest in pursuing advanced uses of networks and I think we ought to push the frontiers. That is clearly essential along with a healthy basic research program.

We need to find a way to involve the communications industry. This is a problem I have been working on directly over the past few years as in seeking a way to allow the graceful evolution of a National Research Net. As long as there is an adequate incentive, I think industry will get involved. If there is no incentive, not only will they not get involved but they could raise serious questions about the process as you indicated.

The government must take the lead in creating the most advanced networks needed by the research community. They simply will not happen any other way. The multi-gigabit per second fiber optic network which you proposed is obviously a good example of an advanced network capability. I believe industry can contribute to that and is prepared to.

Once created, however, such an advanced network will generally be the creature of the Federal Government. That has been our history to date. And, generally speaking, it is provided to the research community as a free good.

Now, if we can find an incentive to have industry provide this capability at some point in the future on a commercial basis, then I think the fruits of this research investment could become more widely available as part of a larger and more powerful domestic information infrastructure. But, in order to do this, we need to find a way to put some element of choice in the hands of the research community.

I think this element of choice is a very key ingredient. Dr. Kleinrock mentioned it and it was clearly discussed in the National Academy Report. This element of choice will allow the research community to exercise its own prerogative as it now does in the selection of computer equipment from multiple vendors.

This situation is not generally present in the more advanced kinds of network environments. When the government provides networking services over the long haul as a free good, the users on those networks have no incentive to seek any alternative source as long as the network is functioning.

But even if industry were interested in getting involved at that level and more advanced levels of networking, the only set of users that exist at that moment are those supported by the government. Even if industry offered equivalent or slightly better service, they would get no takers because they are competing with a free good.

So by putting choice in the hands of the researchers, we can induce some fundamental change in that element of the equation and by "choice," I really mean the ability to spend networking dollars on any equivalent network at the users discretion.

This will have two distinct effects. One, it will provide an attraction to industry to get involved because they will see a market with money to spend. It may be Federal money that is being spent but at least it will provide a financial incentive to industry. Second, it will provide a graceful way for the government to stop supporting special purpose networks which at one time may have been unique and essential but which no longer are, if there is an alternative cost effective commercial choice. There is no good mechanism for doing that today.

Senator **Gore**. With a few possible exceptions.

Dr. **Kahn**. Which exceptions do you have in mind, Senator?

Senator **Gore**. Dr. Huray entered a cautionary word there. He said that there will probably always be some specialized networks. I do not want to endorse that conclusion but I respect his analysis enough to add a caveat that it may well be that some specialized networks will continue.

Dr. **Kahn**. I agree with that. I am saying that once the special purpose nets are no longer unique and essential, then we ought to find a way for the government to get out of supporting them. I believe there will always be some that remain essential and it probably is the role of the government to keep supporting them.

In the absence of such a mechanism, the only alternatives we have are either to maintain the status quo or to actually cause some disruption in the research community as we drop one system and move to another.

In fact, the government often tries to support an old net while it is creating a new one, and that creates financial pressures, generally is costly and does not transfer the old technology effectively into the commercial sector, but rather just lets it die.

I think for a National Research Net to continue to function effectively as a communications subsystem, it really needs conscientious attention to the issue of network management; both technically and administratively. It has to be managed as an entity that needs to be up, working and viable and I think it is really important to recognize this aspect explicitly right from the beginning in such an initiative.

In addition to the normal tasks of management, the day-to-day kinds of operational and planning tasks, there are issues of architecture, standards, graceful evolution that must be addressed, as well as the whole issue of liaison between the research community and industry.

Providing a research net will not necessarily make it useful, as I indicated before. I was involved in this issue at an earlier stage of my career, in the case of the ARPANET which was the pioneering computer network back in the late 1960's. I spent the better part of a year of my life, after the network was up and running and moving bits reliably, working with the research community and the vendors to make the net become usable at the computer level.

We actually ended up planning and holding a network demonstration in the Washington, D.C. area which had the effect of forcing the network into a working status that was then visible to all. I think we are going to need to do something along the same lines as we begin to contemplate a national research network.

More systematically, I think that we need to begin now to develop what I call the higher layers of the research network; or the higher layers of this national information infrastructure. We are talking here about a research network that can move bits from one computer to another. That is but the lowest level of the network. Ultimately, if we can create these higher levels, we will not only improve the productivity of the research community but we will have created a very effective model that could serve as a very useful transfer mechanism into virtually every other sector of society that could use networks.

I think that the highway analogy that you raised is an extremely important one. I have often felt that is a very appropriate analogy. We need to ensure that the highway becomes a useful facility that is more than just laying down concrete. We need the connecting roads, the highway services, and the transportation functions. We are dealing today with a situation in the computer industry that I think is very analogous to the state of automobile transportation at the turn of the century.

We did not have the highway infrastructure then that made it possible for the automobile and transportation to really evolve. Similarly, I think most people cannot do a fraction of what is inherently possible to do with computers today due to the lack of this informational infrastructure.

I will give you two of my favorite examples from a longer list. One is the creation of a digital library system that could be made available on a national basis for the research community. It might allow us to bring together the scientific reports and papers and some of the published research results which are so essential in this community so that they can be straightforwardly accessed by a researcher from any workstation.

You need to know so many things to be able to do that today. It is primarily a technical exercise. It would be impossible to sit down at that workstation and identify what you want without knowing where it is, and how to identify it. It is technically achievable to make such retrievals easy to accomplish if we have the will. However, it ought to be done in a systematic way. By creating a layer of infrastructure above the pure moving of bits back and forth between machines.

A second favorite of mine is a much longer term investment for the country. It would be the creation of a national knowledge bank for Science and Engineering. I envision here starting the process by which we codify the knowledge that is in textbooks and in the heads of scientists and engineers and get it into machine usable form so that some of our more powerful machines, whether they be supercomputers or the artificial intelligence systems of the future, the multiprocessors, whatever, can begin to make use of that. It can be used in flexible manufacturing, for custom design, but I think the most important application in the long term is going to be for education itself. And I think we need to start that process now.

And, finally, I would simply point out that the National Research Net can be a major step forward in achieving collaboration between different research groups. And I think that was pointed to earlier but I think that is a major long-term potential for such nets.

I think we need a national commitment to make such an infrastructure happen and I think that the idea of developing it within the research community is exactly the right way to proceed. We can learn how to make it useful, to find out what its capabilities are. If we do it broadly enough and take the long-term view, this may very well be a precursor for what could be a critical national initiative over the course of the next decade and well into the next century. The National Research Net is an important first step in that direction. I would like to thank you for inviting me to testify and I will be happy to take any questions.

[The statement follows:]

STATEMENT OF DR. ROBERT E. KAHN, PRESIDENT, CORPORATION FOR NATIONAL RESEARCH INITIATIVES

The creation of a National Network for Research is an idea whose time has come. It is needed by the research community and ultimately will be important to society as a whole. It is not often the case that we invest as a nation in specific technological projects to meet our long-range needs and requirements. It is usually the short-term needs which gain the attention and support of key decisionmakers. Often we strive to "get our technology horses back in the barn" after they have departed; yet we may not amply protect, through adequate investment, those still in the barn. Vision and leadership can rectify this. A National Research Network is an example of such vision and leadership and the role of the Government is essential to make it happen.

The OSTP report "A Research and Development Strategy for High Performance Computing" was notable in that it outlined a broad program of Research and Development which responded to more than was originally requested by the Congress. In addition to the research network recommendations, it outlined a broad program of support for High Performance Computing (including supercomputer applications), software, basic computer science research and human resources. In the networking area, it outlined a three-phase initiative to develop a computer network for the entire research community in the U.S., by building on and augmenting the original concept of a fiber optic network for supercomputers. The reasons for this, I believe, are significant. Computer networks and high performance computing have important ramifications for all sectors of the scientific research community. If carried out effectively, a program which addresses these concerns can help maintain a vigorous and highly productive technology base within the U.S.

The three-phase approach to the research network seems very sensible to me. Much of what is outlined for the first phase of the network is already underway and substantial pieces of the network are already operating at 1.544 Megabits/second. For example, the NSFNET backbone was recently upgraded to that speed. The second phase, which would upgrade the network to 45 Megabits/second and stress switching technologies, is within the capability of industry to provide. The third phase, which entails development of a fiber optic net requires considerable research on interfaces and protocols, switching and advanced software applications that can utilize speeds above a billion bits/second. Creating an experimental fiber optic network for a subset of the research community (including some supercomputers) would help to develop the technology for phase three and identify some of the more promising applications for such a network. In my opinion, the most powerful applications will only emerge after the network is available for actual use.

The second and third phases of the plan are currently being considered on a Government wide basis though the mechanism of the Federal Coordinating Council on

Science Engineering and Technology (FCCSET) which is chaired by the President's Science Advisor. Adequate funding for these phases does not yet seem to be in place however. Thus these efforts are necessarily limited in scope and utility. The cooperation of all the key research funding agencies in creating this network is desirable to ensure a single inter-operable network with the opportunity to share resources.

Many basic questions about a research network still remain to be answered. First, how much capacity is required for the various applications of a research network and what can be afforded? Second, what new capabilities will be enabled by the research network? Third, what will it take to create it and how can we leverage the private sector capabilities most effectively in the process?

The Research Network as described in the OSTP report is but the bottom layer of an information infrastructure for the research community. However, it is the foundation on which the other layers can be built. Examples of next layers are (1) A Digital Library System which greatly simplifies and enables access to information from printable documentation to complex data structures from any workstation on the network. (2) A Knowledge Bank for Science and Technology which incorporates textbook level knowledge in computer-based form for use in artificial intelligence applications and education. (3) An Electronic Transaction Framework in which new ideas in computer based interactions may be explored such as flexible manufacturing, rapid procurement, custom design and open bidding.

Careful attention must also be paid to the computational infrastructure which is available to the researchers. In order to make effective use of the research network the research community will require (1) Adequate support for and access to workstations, supercomputers and other intermediate-sized machines. (2) A common evolvable software environment which is conducive to computer mediated interactions. (3) Local area communications and campus networks to link these machines and their software with regional and wide area networks. (4) Higher level software which enables user programs to interact with each other and which also implements the added layers of the infrastructure. (5) Adequate support for (a) research on the use of networks as well as (b) basic research which makes use of these networks.

Other possible applications of a very high speed network include remote access to supercomputers, real-time communication of simulations and graphics, movement of large databases, collaboration between individuals and research groups and telescience. While the actual bandwith requirements for each of these applications will vary with the specific nature of each task, opportunities for experimentation at speeds of 45 Megabits/second to more than a billion bits/second are plentiful.

While the OSTP report addresses certain of these issues, it does not explicitly address the equipment needs of the research community nor the higher layers of the networking infrastructure. The equipment needs (e.g., workstations, gateways, multiprocessors, supercomputers) have traditionally been funded as part of basic research and development efforts. With the advent of a nationwide network for research, the costs of these elements will become large in their own right—perhaps on the same scale as the network itself. The information infrastructure needs for the scientific community have never before been addressed on a national scale.

A National Research Network can bring together these various interests into one coherent experimental system for the research community. This, in turn, can serve as a model for applications in other sectors of society. We should not plan only the lowest layer of networking. We should begin now to plan the higher layers of the research network, even if not all of them can be developed at the same time.

Answers to the following questions will be critical to the success of the network:

(1) How can the participation of all researchers in the U.S. who wish to make use of the research network be assured?

(2) How can the research network be structured to allow cross-communications with the non-research community where appropriate? The analogy to interconnected telephone systems seems appropriate here.

(3) What specific roles should the Government and industry play in transferring the key ideas from the research network into practice for others in society?

(4) How can an orderly process be assured for transition of Government network services into the private sector when sufficient interest develops to make this happen?

(5) How can commercial networking potential be maximally leveraged by the Government to avoid recreating what industry would otherwise be prepared to do with sufficient incentive?

A key to leveraging the capabilities of industry is to put the choice of which networks to use in the hands of the individual research managers. Today, for high speed networks, these choices are largely made for the users by the Government. The user receives, in effect, network service as a free good with little incentive to use the network resources efficiently. Nor is there incentive to seek other providers and thus no incentive for industry to get involved.

The Government must take the lead in creating the most advanced networking capabilities that technology will alllow. These "seed networks" such as the proposed phase three fiber optic network, will allow the technology to be created by industry and will allow the research community to learn how to use it. Over time these "seed networks" will mature as the user community grows. Ultimately, industry should provide these services commercially.

How can we ensure that this process moves at a steady pace? If there is but a single advanced network supported by the Government with a given class of service (such as an experimental fiber optic network running at a speed of a billion bits per second) then the Government may as well decide on behalf of the users who will get access. There is no alternative to choose from. When one or more commercial networks appear with equivalent service the users should be allowed to decide which one to use. The ability of a user to choose may be the single most important factor in the decision by industry to develop and provide such equivalent services. The exercise of choice may have the effect of permitting the original Government network to be downsized if the commercial alternatives are attractive enough technically and economically.

Several ideas have been suggested as to how this exercise of choice may be achieved. No matter what technique is selected, however, the individual departments and agencies of the Government should have the final say in allocating Government resources to research contracts. This could be done by providing the equivalent of fungible network vouchers or subscriptions. These vouchers could then be applied to any qualified and equivalent commercial or Government network service.

I strongly support the concept of a National Research Network and urge the Congress to move forward on this issue. The approach outlined above will provide flexibility for the Government in transition from Government funded networks to commercial services which best meet the continuing needs of the research community. In your deliberations, I hope you will leave room for the network concept to grow and flourish along the various dimensions I have outlined.

[The following information was subsequently received for the record:]

QUESTIONS OF THE CHAIRMAN AND THE ANSWERS THERETO

Q: You've mentioned that advanced networks will give researchers access to powerful data bases that could help keep America competitive.

- Should we be worried that those same data bases might be used by other countries to compete against our industry?
- How easy is it to protect a data base on a network from unauthorized access?
- What techniques are there for ensuring that no one has tampered with the data in a data base?

A: I would distinguish between basic scientific databases which might typically be accessible on a future national research network and certain critical databases which would be protected via administrative and technical measures. U.S. scientific productivity has been high because we have encouraged ease of access to scientific information. Countries that have adopted closed policies in this area have usually fallen behind. While maintaining an open environment for the conduct of scientific research is essential, we must also encourage and develop better methods for coupling results of this research to industry. In addition, increased attention at a national level to the competitive situation is warranted to monitor progress abroad and so that we may respond, as appropriate.

There are many ways to protect database information, with varying degrees of security, as there are many ways to protect individual property. Participants in advanced networks should be free to choose the level of security that best meets their needs. Of course, the ultimate in protection occurs when access is completely closed, but such steps would clearly be counterproductive for the scientific community as a whole. Administrative procedures, such as passwords, are sufficient to deter most ordinary forms of unauthorized access. Encryption can be used to protect data in transit. However, no system of protection is completely foolproof and human beings can subvert an otherwise well protected system. Legal remedies are the only recourse when human systems fail. I think tampering is not likely to be a significant problem in a national research network, which offers quite a different context than, say, banking where the issue would be primary. When the early work has been accomplished to establish a national research network backbone, techniques such as "audit trails", screening of entries by trusted individuals, and a variety of cross-checking mechanisms can be used in those specialized database situations where the need for such steps may be warranted. For most situations involving the research community, there should be no real cause for concern here.

Q: There have been dozens of articles in the press lately about the threat of computer viruses. It seems to me that a national research network will help spread these viruses.

- How serious a threat are computer viruses?
- Are there effective ways to protect computers from this threat?

A: Computer viruses and their close relatives (e.g., worms, trojan horses) are a potential threat to the research community. However, such programs can enter a machine with or without computer networks and carry out functions which were not intended by the user or the operator of the system. I see these threats primarily as sociological threats. The telephone network can transmit viruses between computers as can networks specifically intended to connect computers. Individuals can also propagate virus programs by simply carrying storage media from one machine to another. In at least one case, it has already been reported that original manufacturer software was infected at the source and distributed commerically without the knowledge of the manufacturer. This problem was quickly corrected as was the recently publicized worm. It is becoming increasingly difficult to find surreptitious ways to infect current computer technology. Yet, occasional disruptions will occur due to such misuse. Again, in my opinion, the benefits of an open national research network outweigh the occasional risks of misuse. However, research into computer and network security should continue to be pursued vigorously to learn how to build more robust systems and to further protect against possible misuse in the future.

Q: The OSTP report calls for developing networks that can transmit billions of bits of data every second. If the network runs 24 hours a day, 365 days a year, that adds up to a lot of data.

- How do we decide when we've got enough capacity?
- Although a billion bits a second is an impressive number, if you divide it by a million researchers, you don't have much capacity per researcher. Will the projected increase in the number of network users keep us from seeing the benefits of increasing its capacity?

A: Here we set aside the issues of cost, and focus only on the utility of multi-gigabit/second networks. An individual user will require billions of bits per second for real-time transmission of high resolution images, simulations and 3-D graphics. Additional capacity will be required to support simultaneous use of the network by multiple researchers in order to avoid diluting the benefits of very high speed transmission. However, not all researchers will require such high data rates in the same sense that not all researchers require use of the fastest supercomputers. Further, the available capacity, which consists of multiple paths through the network will be shared among many researchers of which only a portion will be active on the network at any given time. Like the telephone network, the statistics of use of a national research net will likely exhibit a "locality of demand" in which the bulk of a researcher's capacity needs should be satisfied within the boundaries of his research site. This means that the capacity of the national research network backbone is needed to serve only a fraction of total user demand at any particular time.

By exploring the limits of technology in our nation's research laboratories we stimulate and hasten its productive use in commercial applications. Quantitative changes in technology often lead to qualitative changes in its application. Such advanced research has lead to an advantage for the U.S. in many other cases (e.g. supercomputers). I believe similar opportunities will occur when a gigabit/second network is available for use by the research community.

Q: In your testimony you've mentioned a number of interesting applications for a national computer network, including a digital library, an electronic transaction framework, a medical information exchange, scientific data banks, and facilities for custom manufacturing. Do you have any estimates of the cost of developing and running these different projects?

A: The Corporation for National Research Initiatives has proposed a concept of operations for a national digital library system to meet the needs of the scientific research community. This is an advanced network-based information deposit and retrieval system that can assist a user in storing and retrieving information and in obtaining transparent access to a variety of on-line services. Estimates for the research and development to create such a system are $20-30 million dollars per year for a 5-10 year period. The current thrust of the proposal is on developing powerful mechanisms for simplifying and controlling access to remote databases and creating effective user interfaces with full accountability. Once developed, the cost of operation will depend on the total number of active users and the total number of databases supported by the system. It is reasonable to assume this cost will initially be approximately equal to the annual cost of research and development on the system. This assumes the use of existing databases. Funds for development of specific databases are not included in this figure. Estimates for the other projects are not available at this time. However, my personal experience in developing and managing systems of the magnitude described leads me to conclude that an annual expenditure for network applications on the order of $100 million - $1 Billion per year would be appropriate. If we spend several hundred million dollars on a national research network, we should support applications at an adequate level to make it genuinely useful to the entire scientific community.

Q: You've emphasized the need to give researchers a choice of networking options.

- Could you explain further why this freedom of choice is so important?
- What mechanisms might we use to ensure that researchers have a choice?

A: There is a "chicken and egg" problem in getting industry to develop advanced network technology and services. They will not invest until they can see a market of significant size developing within a few years. On the other hand, a market may not develop until the network technology and services come into being, so that the research community can explore them. The idea here is for the government to take responsibility for funding the most advanced research and development and to provide research network capabilities until such time as they can be obtained from industry. The national research network consists of commercial networks, private networks and government networks. Industry cannot justify creating network capabilities which are only required by the research community initially and which are otherwise provided free by the government. However, if the users of the government networks are given a choice, they may elect to use a commercial alternative if it offers better service at a lower price. This expression of choice allows a migration path from advanced networks developed and operated as one-of-a-kind resources by the government to commercial systems.

There are several mechanisms one might contemplate. The one I favor is some kind of voucher system which is used to provide government contractors access to a network of their choice (within a particular class). If each government network has a designated fee structure, and the vouchers are fungible (i.e. convertible by the networks into money), contractors can spend the vouchers on commercial networks or on government networks. This provides the needed incentive for industry since it gives them access to the only near-term market that exists for that network service - one that was previously inaccessible. If such a voucher system is established, a transition mechanism should be provided to support the government networks while they are still in widespread use but losing customers to commerical systems. This can be accomplished in a way that insures that government networks are downsized in a stable way until their services are deemed by the government to be no longer necessary. Of course, the government must retain its role in fostering the continued development of the most advanced networks as the older maturing networks transition.

Senator **Gore**. Thank you very much. We will hold questions until the panel has completed testimony. Dr. Robert Dickinson is Deputy Head of the Interdisciplinary Climate Systems section at the National Center for Atmospheric Research.

Welcome, please proceed.

STATEMENT OF DR. ROBERT E. DICKINSON, DEPUTY HEAD, INTERDISCIPLINARY CLIMATE SYSTEMS SECTION, NATIONAL CENTER FOR ATMOSPHERIC RESEARCH

Dr. **Dickinson**. Thank you, Senator Gore. I am here as a scientific worker who uses supercomputers and to a certain extent networks. I would first of all like to thank you very much for your scientific leadership in this and other scientific areas. I believe in my community we already think of you as an honorary member.

Supercomputers are essential for our research on climate and weather. And they represent one of the major tools for projecting future climate from increasing trace gases, as one particular example that is of concern.

We now can talk in using these models about global warming and we understand well enough what is going on scientifically to raise this as a very important concern but we are not currently able to use these models down to the level of regional climates where everybody would like to know the answers.

Thus I project that because of demand for more knowledge as to what will happen in the future, models of climate systems are going to need much more accurate simulations which will require increasing computing capacity of many factors of ten. I said a thousand in my written testimony; some large increase anyway will be required over the next decade or two.

The need for these capabilities may be met by faster and more numerous computers and possibly also by facilities that are more focused on particular priority questions such as the one I have raised here.

For climate studies, besides the computers, we also will have to develop in parallel better observational data sets and increased understanding of the system we are looking at. Certainly in terms of increased understanding, communication is a very important aspect and networks do and potentially can play a much larger role in that area.

Computers that we have now and in the future must be made to be used by a wide variety of researchers and for this they must be easy to use by scientists without very much special training beyond their scientific knowledge. These criteria then require that the computer community have a considerable investment in software research and development, such as has already been addressed at this hearing, for developing the new technologies that facilitate using these massively parallel processors.

Right now I think I could speak for all my colleagues in saying that we are a bit leery about using future systems that looks very complicated to us to run the sorts of problems we now have in computers that we understand how to use. We certainly endorse the idea of having a balanced approach to not only hardware of computers but software and networks.

Networking is a major asset to increase the efficiency of our scientific computations and our communication. Its linkages enhance our focus on important questions and reduce duplication of efforts. Our community should all be linking together and sharing its information because we cannot continually have everybody reinventing the wheel.

A personal example of this is a scientific visitor from Berlin who came to NCAR last spring and used a big software package I had developed. He went through the code line by line, found a few minor bugs, suggested some improvements and then took it home with him back to Berlin for some uses of his involving satellite remote sensing. Well, this sort of thing could be done much more efficiently and rapidly by networks where one could talk with colleagues and exchange the software through such networks.

In terms of your earlier question on the global connections of networks, this is very important to us, as illustrated by my example. The problems we work on are global and it is important to be able to extend what we do to colleagues in other countries, both from the viewpoint of efficiency, but also from the viewpoint of having any hope of having something done about global environmental problems. Because they are global, there has to be a global consensus that is reached and this can only be achieved by having the global community participate in research. I guess that will wrap up my remarks.

[The statement follows:]

STATEMENT OF ROBERT E. DICKINSON, NATIONAL CENTER FOR ATMOSPHERIC RESEARCH

1. PERSONAL BACKGROUND

I speak as a scientific user of supercomputers and networks. My research uses complex general circulation computer models to understand the dynamics of the climate system and how it might change. I am involved in improving our capabilities for modeling the land component of the climate system, including the role of vegetation.

Our development of general circulation models (GCMs) in NCAR's Climate and Global Dynamics Division involves approximately a dozen Ph.D. scientists, a comparable number of support staff, and a large body of university collaborators. Even with this level of effort, our tasks at times can seem overwhelming. The number of disciplines we must master well enough to include in our models is considerably greater than the number of scientists constructing the models. Our models use supercomputers (currently a CRAY X-MP) provided by NCAR's Scientific Computing Division. My technical expertise in supercomputers and networking is fairly limited. My background deals primarily with the scientific questions requiring these technologies.

I have included in the references several reviews that have recently addressed the issues of modeling long-term climate change. A considerable fraction of my recent time has been spent in community research planning activities through various committees of the National Research Council (NRC) of the World Climate Research Programme (WCRP), of The International Geosphere-Biosphere Programme, and of the American Meteorological Society (AMS) and American Geophysical Union (AGU). I am chairman of the Climate Research Committee of the NRC, a member of the U.S. Global Change Committee, and a member of the U.S. National Academy of Sciences. My testimony is personal and does not necessarily represent the views of any of these organizations nor of NCAR.

2. CRITICAL SCIENTIFIC APPLICATIONS

The research communities I work with share the concern of this committee for potential threats to the global environment imposed by large-scale human activities. I have been involved in the issues of global climate change from trace gases and impacts of tropical deforestation and ozone-layer chemistry. The details of these issues are not the point of this testimony. Rather, I address the question of improvements in some of our research tools important to the accuracy of our models.

Computer modeling is used by my community to achieve as accurate as possible a description of the future global environment given our current understanding, computational resources, and scenarios for past and future societal activities. Better models, and hence reduced uncertainties of the future, require improvements in three areas. First, we must expand the resolution of our models, which means larger streams of numbers must be calculated and moved around. Resolution is the number of points over the earth treated by the model, hence the degree of freedom of the model. Bigger, faster, and more numerous (i.e., cheaper) computers are the main means to achieve such higher model resolution. Second, we must improve the mathematical descriptions of physical processes represented in our models. The primary limitation of progress is the shortage of skilled scientists in this area and the quality of their communication. Third, we need intensive field programs and global data sets to better test and validate the individual processes within our models. Many international programs we participate in through the WCRP are responding to this need; the largest planned future effort is NASA's Earth Observing System (Eos) program. The analysis and utilization of data from process studies and of global data sets also requires distributed and central computing facilities and communication systems.

Our models for global climate research treat the same dynamical system as do numerical weather prediction models, so not surprisingly, we share the same basic model structures and history of computing advances. Climate research, however, emphasized more heavily the physical processes that determine the system behavior on time scales of years and longer, e.g., land and ocean surfaces, clouds and radiation, snow and ice, soil moisture, etc.

3. REQUIREMENTS OF ATMOSPHERIC MODELS FOR UPGRADED COMPUTING FACILITIES

Atmospheric models use supercomputing resources in two distinct modes: operations and research. The development of supercomputer technologies continues to be a pacing item for progress in both areas. Operational activities require specific calculations using a single model or a small number of models. This mode is currently restricted to weather prediction centers, e.g. the U.S. National meteorological Center of NOAA and the European Centre for Medium Range Weather Forecasts (ECMWF). At the other extreme is NCAR where several hundred in-house and university scientists share the computational resources for exploring a wide range of scientific questions. Because operational centers have more focused modeling activities, they have only a few computer programs that they run regularly. This is contrasted with the many thousands that go through the NCAR facility. Thus, operational centers can more readily tailor their calculations to their present computer and demand computational output within some time limits, e.g., ECMWF requires 10 day forecasts to be produced within 5 hours.

One view of climate is that it is some kind of average of all the day-to-day weather events. Weather forecast centers, in particular ECMWF, have demonstrated that the weather systems simulated by global models continue to become more realistic with increasing horizontal spatial resolution down to grid spacings under 100-km. Research with regional mesoscale models also demonstrates this point. Yet, computer resources limit much of my research to models with 500-km or so resolution. Figures 1 and 2 illustrate the differences in grid-point distribution and in topography over the western United States for a 500-km versus 60-km resolution. As a Coloradoan, I doubt whether one grid point and a smooth Great Plains terrain will represent my climate, past or future. I doubt that Nevadans seem themselves on the western slope of the rocky Mountains and represented by an average of a California and a Utah grid point. The figure shows what you get for these states with present global models. Likewise, few Eastern states get a single global model grid point. I believe more that a factor-of-10 increase in resolution (requiring up to a thousand times more computation because this increase corresponds to a factor-of-100 increase in number of grid points, and shorter time steps are necessary) is a prerequisite to adequate modeling of the regional aspects of global climate change. With regard to future climate changes, I believe that present models are good enough to suggest possible futures, but without the increased resolution they will not be good enough to determine what will happen at the level of detail we would like.

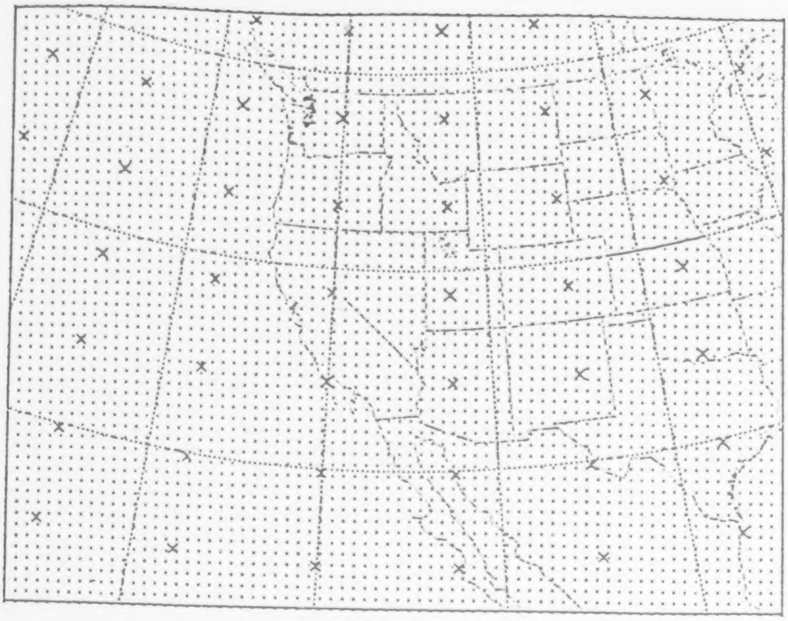

Fig. 1 Distribution of mesh points over Western United States for NCAR's current global model (the x's) and for its mesoscale model (the dots giving a 60 x 60 km mesh).

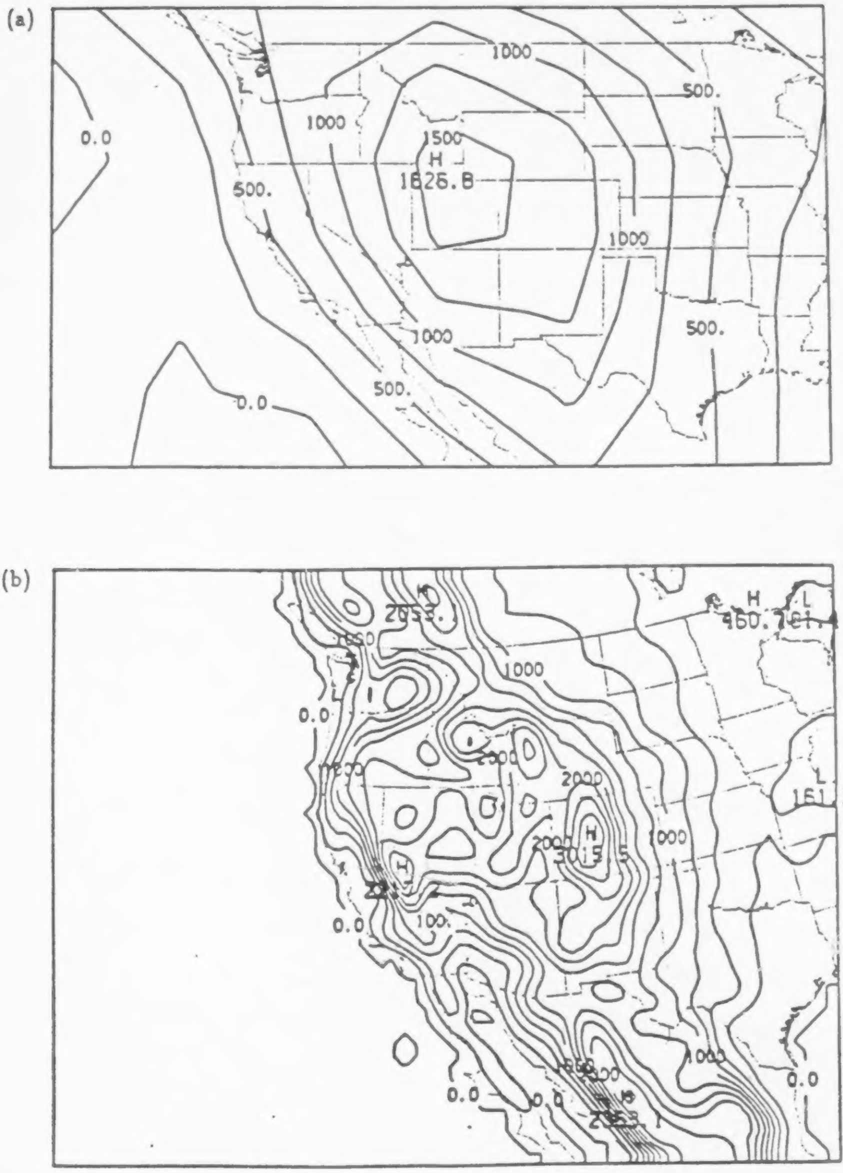
(a)
1000
500.
0.0
1000
1500
H
1626.8
500.
1000
1000
500.
500.
0.0
(b)
H
2053.
H
L
1000
0.0
2000
2000
H
3015.5
1000
L
161.2
0.0
1000
0.0
0.0
0.0

Fig. 2 The mountains of western United States given in meters elevation (1 meter = 3.1 feet), with contours every 250 meters, are used for a) the global model (which must be highly smoothed to match the model resolution, and b) the 60 x 60 km mesoscale model (which while still smoothed, distinguishes between e.g., the Sierras and the Rockies).

Nationally, we have seen approximately a factor-of-30 increase in computer power every 10 years. At my institution, the resource per scientist has grown more slowly, perhaps an average factor-of-5 increase per decade and even less for the largest users because of the corresponding large increase in the size of the community using the resources. The above requirement of a factor-of-1000 increase in computer resources does not necessarily suggest the need for an equivalent speedup in the most advanced computers for research purposes, in contract to those for operational computations. It could, in principle, be partially satisfied by large increases in the number of computers with capabilities comparable to present supercomputers. Larger models require not only greater effectiveness of central processor units, but also larger memories and mass-storage facilities. Something like 1% of the numbers generated by our models is saved on a mass-storage system for postprocessing of model calculations, and these represent some of the world's largest data sets. The capabilities of the mass-storage system determine the quality of the information that can be extracted from the model output.

Scientists work much more efficiently if they can complete individual model calculations within a few days. Ocean and mesoscale modelers at NCAR have comparable requirements for improvements in their capabilities.

Because of many other demands, few of our scientists have much time to devote to the special programming requirements of multiprocessing computers which could shorten the time required to complete their calculations. A few models at NCAR use the multitasking facilities of the CRAY-XMP, but most scientists use a single individual processor, i.e., they use the computer as an equivalent to 4 individual processors. Unless considerable software research and development facilitates their use, it may be at least a decade before the software for massive multiprocessing will be far enough advanced to be usable by most of our scientists. Thus, while "faster and more" would appear better, some of the newer hardware technologies may not be immediately usable by us without additional programming and software support.

Having indicated a considerable requirement for major increases in computational resources for atmospheric sciences and climate modeling, I repeat earlier caveats. Our scientific understanding and manpower, our observational data sets and their analyses, and our software must be correspondingly improved to reap optimal benefits from such resources.

4. BENEFITS OF IMPROVED NETWORKING

Networks have been replacing less efficient procedures for doing similar functions. By facilitating the computations and scientific communication, they help immensely in enhancing our limited research manpower. For example, all of our computing traffic at NCAR, now carried through a local network, previously used IBM-punched cards submitted to operators (I still have a few such decks in my office.) The current procedures are somewhere between two and ten times as efficient in use of scientists' time.

University scientists used to spend their summers at NCAR computing. Now they use our facilities year-round from their home institutions through networks. I can't speak for them, but we benefit from a less crowded building although we lose some direct communication with them. These networks involve transfers of computer programs and small amounts of program output. Large amounts of output, e.g., data sets generated by climate simulations are not usually transferred by the present relatively slow networks. Rather, they are more often put on 6250 BPI physical tapes and carried within the building or mailed to offsite users. The latter method takes at least several days. How important it is to use networks depends on the eagerness of the offsite scientist and relative costs. Again, immediate access to the model output could increase productivity considerably. As we employ modern visualization techniques (color and animated graphics) to look at model output, high-speed networks will become yet more useful.

An increasing part of our scientific lives is network communication. I make some use of the INTERNET computer-to-computer system, but in the last year have seen considerably more such mail through commercial telephone line systems, e.g., TELEX, FAX, and the Telnet system, i.e., OMNET AGU's KOSMOS, and NASA-MAIL. All of these facilities together increase the speed and reach of scientific communication. Ultimately, computer networks could be the most powerful form of communication because of their ready exchange of communications and software. But they must be more general, more widely available, and easier to use.

5. CONCLUSIONS AND SUMMARY

Supercomputers are essential for climate and weather research, as well as for research on other scientific questions.

Increases in computational capabilities by up to a factor of a thousand, will be needed within the next decade for modeling important scientific questions, such as accurate simulations of regional climate change within the framework global models.

The need for these capabilities may be met by faster and more numerous computers and by facilities more focussed on priority questions.

For climate studies, better observational data sets and increased understanding are also high priorities.

Computers for application to a wide variety of research must be easy to use by scientists without special training. These criteria require a considerable investment in software research and development for the new technologies using large numbers of central processors.

Networking is a major asset to increase the efficiency of scientific computations and communication. Its linkages enhance our focus on important questions and reduce duplication of efforts.

Senator GORE. Thank you. I will hold off on questions.

Our next witness is Dr. Paul Kutler, Acting Director of the Aerodynamics Division and Office of Aeronautics and Space Technology.

Dr. Kutler, welcome, please proceed.

STATEMENT OF DR. PAUL KUTLER, ACTING DIRECTOR, AERODYNAMICS DIVISION, OFFICE OF AERONAUTICS AND SPACE TECHNOLOGY, NATIONAL AERONAUTICS AND SPACE ADMINISTRATION

Dr. KUTLER. It is a pleasure to be here. I would like to talk a little on networks and applications. I welcome this opportunity to appear before you to discuss NASA's effective use of supercomputers in aeronautics and space and a continuing requirement for state-of-the-art computational facilities, including network subsystems for enhancing this country's competitive posture. First chart, please.

Today's supercomputers are used by NASA scientists in a variety of disciplines. In my testimony today I will emphasize one of those. One of the most dominant and fastest growing uses of supercomputers today is aerodynamic simulation. In partnership and experimentation and flight tests as you see on this chart, computational fluid dynamics, CFD as we like to refer to it, provides an additional simulation tool that possesses significant advantages including physical realism and cost-effectiveness.

CFD for the aerospace industry has basically two objectives, to reduce the development cost and increase the performance of aerospace vehicles and/or their components.

The discipline of CFD involves the use of supercomputers for the numerical solution of mathematical equations that describe the behavior of fluids about those aerospace vehicles and their components.

Next chart, please.

An example in which CFD was used to model flow over a complex configuration is the integrated space shuttle launch configuration. This was a collaborative program between NASA's Ames Research Center, Johnson Space Center, and Rockwell International Corporation.

The objectives of the simulation are to understand the flight performance of the launch configuration and use the results to explore the feasibility of crew escape during launch.

Computer graphics plays an important part in the reduction of data for these complex configurations and I have a three-dimensional viewer here that depicts flow about a space shuttle launch configuration and shows the pressure about that configuration, comparing both the experimental data on one side and the numerical data on the other side. You will see this same picture in the video we plan to show in a few minutes.

Numerical simulation has also been applied to internal flow problems, an excellent example of this is the effort to simulate the flow and components of the space shuttle main engine with the objective of increased engine performance and lifetime. In this effort the flow is simulated in the SSME or space shuttle main engine hot-gas manifold. It is a collaborative program involving NASA's Ames Research Center, the Marshall Space Flight Center, and Rocketdyne, a division of Rockwell International.

CFD technology has also been developed and exported to aircraft companies that will enable them to model flows about realistic aircraft configurations, one of which will be shown in the video.

I now would like to show you this video. It is about four minutes long and it more vividly depicts the numerical simulations that have been generated. The two examples are about the space shuttle configuration and the main engine and the F-16 fighter aircraft.

[A videotape was shown.]

Dr. KUTLER. It is very important to implement, utilize computer graphics in the reduction of the data because we have a three-dimensional data base and in some instances that data base is unsteady flow, so it involves an enormous amount of data. The only way you can digest, understand and make corrections or changes to the configuration is by displaying that data using our graphics display devices.

One of the first steps in modeling the flows about configurations like this is to do the geometry definition. As you saw there, they very accurately modeled the geometry. What you are seeing now is a comparison of the pressure distributions of a launch configuration for a Mach number of 1.55 and an angle at the time of −6 degrees.

The support structure that holds the orbiter to the external tank was not modeled in this case. We have subsequently provided the calculation with that modeling. As you can see, the numerical results compare quite well with the experimental values and help to validate the CFD technology.

We have also applied this to internal flow configurations. And what you are seeing now is a component of the space shuttle main engine called the hot-gas manifold.

The current design on the space shuttle main engine consists of three transfer ducts. The configuration on the right is a two-transfer duct case, configuration that was designed using computational solutions in combination with experiment. As you can see, there is a small variation in colors there which indicates small pressure loss. The two duct configuration is an improved design.

As a result of that program, a more efficient two-transfer duct case was designed with higher power levels and reduced turbine bearing loads.

CFD application software has also been developed to compute the flow of the F-16 aircraft, including the flow into the engine inlet up to the compressor and out of the jet exhaust turbine. It is capable of modeling flows about these configurations at large angles of attack as well as with the aircraft yawed to simulate performance near the flight envelope boundaries. The calculation shown here is for a Mach Number 0.9 at an angle of minus 6 degrees. They require 500,000 grid points to model the flow region.

The numerical solution has been validated against experimental data. A calculation such as this requires about 25 hours of computer time on this country's fastest computer. Calculations as sophisticated as these were enabled by the advanced numerical analysis solved by NASA scientists and the increase in power afforded by the supercomputer.

The computational aeroscience challenge, disciplinary simulations of aerospace vehicles in flight, is combining the disciplines aerodynamics, structures, flow chemistry, and electromagnetics. To accomplish this will require computers capable of 1 trillion floating point operations per second on a sustained basis. That is 1,000 times faster than computers we have available today.

The development of massively parallel computers and the associated software offers a means for satisfying this goal. Internally NASA is currently looking at how to respond to the Office of Science and Technology Policy report entitled, "The Research and Development Strategy for High Performance Computing."

High performance computing necessary to solve NASA's many challenges will require coordinated effort with other Federal agencies that has been pointed out here.

It will also require synergistically combining the expertise of physicists and computer scientists and coordinating a development effort for application software, system software, hardware and the peripherals.

The next chart, please.

Computer systems required by the discipline of CFD are composed of not only the high-speed processor but subsystems such as pre- and post-data processors, on-line high-speed and off-line archival storage devices, a local area network and high-speed long-distance networks.

Effective and efficient long distance networks are critically important to foster synergistic cooperative programs among organizations, two of which we have just talked about.

NASnet of NASA's Numerical Aerodynamic Simulation Facility is a key ingredient of the NAS system and links remote users at other NASA centers, industry sites and selected universities.

A map displaying the existing and planned NASnet configuration is shown in this chart. NASnet consists of nineteen 56 kilobit per second switch circuits, two dedicated 1.5 megabit per second circuits and three 9.6 kilobit per second circuits. In addition, all NASA installations are connected by the Program System Communication Network (PSCN) which provides a backbone for NASnet as shown by the yellow line there.

In summary, NASA computational fluid dynamics program paced by physical modeling and the computing system has played a key role in the aerodynamic simulation process and has impacted the aerospace vehicle design process through cooperative programs with industry.

Such programs require an effective and efficient computer network as part of the computational facility for electronically linking participating organizations. NASnet, of NASA's Numerical Aerodynamic Simulation Facility, is a successful example of a disciplined network that supports cooperation, possesses sufficient data transfer capacity, interconnects effectively with other networks and ensures privacy of use.

Progress in computing can be accelerated through the continued pioneering of hardware, software, algorithms and peripheral technology and the effect of commercialization of that technology.

NASA's growing emphasis on computational aerosciences and supercomputing technology with an objective of developing validated application software today for the computers of tomorrow will help to ensure this country's preeminence in aeronautics.

Thank you, Mr. Gore.

[The statement follows:]

STATEMENT OF PAUL KUTLER, ACTING DIRECTOR, AERODYNAMICS DIVISION, OFFICE OF AERONAUTICS AND SPACE TECHNOLOGY, NASA

Mr. Chairman and Members of the Subcommittee:

I welcome this opportunity to appear before you to discuss NASA's effective use of supercomputers in aeronautics and space and the continuing requirement for state-of-the-art computational facilities including network subsystems for enhancing this country's competitive posture via cooperative programs with industry for rapid technology transfer and implementation. NASA has enjoyed a good working relationship with industry and has readily shared its scientific expertise, technology and facilities to enhance this country's aerospace products. In the future, we anticipate this synergistic relationship to become stronger being spurred on by the growing capabilities of international competitors.

Today's supercomputers are used by NASA scientists in a variety of disciplines including structures, propulsion, chemistry and earth and space sciences to name a few. One of the most dominant and fastest growing uses of supercomputers, however, is aerodynamic simulation; this discipline will be emphasized in today's testimony. In partnership with experimentation and flight testing, computational fluid dynamics (CFD) provides an additional simulation tool that possesses significant advantages including physical realism and cost effectiveness. In some instances it is the only simulation tool available to the aerospace vehicle designer. For example, it is an enabling technology in the National Aero-Space Plane program because of the lack of ground-based, hypersonic experimental facilities. Computational fluid dynamics for the aerospace industry has basically two objectives: (1) to reduce the development cost and (2) to increase the performance of aerospace vehicles and their components.

Computational fluid dynamics involves the use of supercomputers for the numerical solution of mathematical equations that describe the behavior of fluids about aerospace vehicles and/or their components. Solutions from CFD are "information rich" and provide the engineer with data requiring high-speed processors and sophisticated graphic workstations for display and analysis. The discipline of CFD is paced by physical modeling (including turbulence and chemical reaction) and the computational system. With increased computer power comes a decreased reliance on physical modeling and, thus, reduced uncertainty in the numerical solutions.

An advantage of CFD over other aerodynamic simulations is that its cost is decreasing with time. This can be attributed to two factors: (1) the decreasing cost of supercomputers coupled with their increasing capabilities and (2) the improvement in numerical procedures for solving the mathematical equations.

Computer systems required by the discipline of CFD are composed of not only the high-speed processor, but also various subsystems such as pre- and post-data processors, on-line high speed and off-line archival storage devices, a local area network and a high-speed long-distance network. The primary, and most costly, element in the system is the high-speed processor. However, for optimal operation, it is necessary that the entire system be balanced and tuned for maximum productivity. Any weak link creates a bottleneck in generating information. NASA's Numerical Aerodynamic Simulation (NAS) Facility is an example of a balanced system that is user responsive. Its charter provides for a continual equipment upgrade to maintain a state-of-the-art facility.

A key ingredient of the NAS system is its network called NASnet that links remote users at other NASA centers, industry sites and selected universities. Effective and efficient long-distance networks are critically important to foster synergistic cooperative programs among these organizations. NASnet consists of nineteen 56 kb/s switched circuits, two dedicated 1.544 mb/s circuits (T1), and three 9.6 kb/s circuits. In addition, all NASA installations are connected by the Program Support Communications Network (PSCN) which provides the backbone for NASnet. A map displaying the existing and planned NASnet configuration is shown in Figure 1 At the present time all NASnet operational circuits are terrestrial because of the relatively poor response time inherent with satellite communications. The hurdles raised by satellite communication are being addressed by NAS personnel.

Numerical simulation has made great strides in the detailed modeling of complex fluid flows, including the fluid physics of flows heretofore not understood. It is routinely being applied now to some rather complicated problems and impacting the design cycle of aerospace flight vehicles and their components. An example is the numerical simulation of the flow over the integrated Space Shuttle launch · configuration. This was a collaborative program involving NASA's Ames Research Center, Johnson Space Center, and Rockwell International Corporation.

The objectives of the simulation are to understand the flight performance of the launch configuration and use the results to explore the feasibility of crew escape during launch. A comparison of computed and experimental results for a Mach number of 1.55 and angle of attack of -6 degrees is shown in Figure 2. Although the support structure between the orbiter and the external tank was not modeled, the computed results compare well with the experimental values and help to validate the CFD technology. A calculation as sophisticated as this was enabled by the advances in numerical methodology developed by NASA scientists and the increase in power afforded by today's supercomputers. Future work on this configuration will extend the external flow simulation to include staging, where the orbiter separates from the solid rocket booster and external tank assembly.

Numerical simulation has also been applied to internal flow problems. Two excellent examples of this include efforts to simulate the flow in components of the Space Shuttle main engine (SSME) with the objective of increased engine performance and lifetime. In the first effort, a collaborative program with NASA's Ames Research Center, Marshall Space Flight Center and Rocketdyne (a Division of Rockwell International), the flow was simulated in the SSME hot-gas manifold. As a result of that program, a simpler, more efficient design (two fuel ducts instead of three), with higher power levels (5 percent increase), increased reliability, and reduced turbine bearing loads (80 percent decrease) resulted. Certification of that engine for flight is in process.

In the second effort, the flow in the SSME high-pressure fuel turbopump, two-stage turbine was simulated numerically by Marshall Space Flight Center using CFD technology developed at the Ames Research Center. The calculations provided engineers with valuable information as to the thermal and aerodynamic loads on the turbine blades that can be used to enhance engine design and provide better estimates of engine life. That same computer code, originally designed for jet-engine turbine analysis, has been transferred to several aircraft engine manufacturers and is being used by them to enhance their designs and reduce engine development cost. An increase of 1 percent in turbine efficiency of all jet engines on America's transport fleet would result in a fuel savings of $100 million a year. This translates into lower ticket prices for the consumer.

Computational physics development, to date, has taken place on single processor computers. With advances in clock-speed and memory size, largely brought about by more efficient chips, the overall processing speed has increased by an order of magnitude during the past ten years. The physical limits of such development are now being approached, caused by inherent electron speeds and problems of cooling. Although new opportunities may be provided by the development of practical means of achieving near-room-temperature superconductivity, the direction most promising for the future is parallel processing.

Supercomputers available today generally have more than one processor and thus are considered lightly parallel. The recently announced Cray YMP has eight processors and is 3-4 times as powerful as its predecessor, the Cray XMP 4/8. Conventional supercomputer technology is likely to increase in parallelism at a moderate rate with steady factors of 3 to 4 improvement in performance over the preceding model. Today, however, the conventional designs are being challenged by massively parallel architectures offering thousands of processors and the potential for an order-of-magnitude performance improvement in the near term over conventional supercomputers. The impediments to the rapid application of massively parallel computers to computational physics problems are the availability of the hardware and the limited amount of system and applications software for those machines. Formulation of computational physics problems suitable for solution on massively parallel processor machines and the internal management of the data among the processors are first-order research problems at the present time.

High performance computing is essential to progress in science and engineering. The lead in high performance computing that the U. S. has enjoyed is being challenged from abroad. NASA has been a leader in computational physics and one of the pioneering agencies in supercomputing facilities. Internally, NASA is currently looking at how to respond to the Office of Science and Technology Policy report entitled, "A Research and Development Strategy for High Performance Computing," to enable the solution of grand challenges raised in the report.
NASA has defined grand challenges (fundamental problems in science or engineering whose solution would be enabled by future high-performance computers) in two disciplines: (1) computational aerosciences and (2) computational Earth and space sciences.

Today's computational aeroscience challenge is to perform integrated, multidisciplinary simulations of aerospace vehicles in flight. This will require simulation of propulsion, external aerodynamics, structures, controls, chemistry, and electro-magnetic phenomena in an integrated and optimized manner. It is currently estimated that this type of calculation will take between 10-100 hours on computer systems capable of sustaining a trillion floating point operations per second (teraflop). Likewise, in space and Earth sciences, NASA has made significant strides in weather/climate simulation, simulation of stellar/galactic evolution, solid Earth modeling, and analysis of multi-spectral/multi-sensor data. Today's space and Earth sciences challenge is the multidisciplinary modeling of the processes and the coupling of those processes from the center of the Earth to outer edges of the universe to better understand the creation of the universe, the origin of life, and the future environment of man. To begin to do this modeling, a teraflop computational capability is required.

High-performance computing necessary to solve NASA's grand challenges will require a coordinated effort with other federal agencies. It will also require synergistically combining the expertise of physicists and computer scientists in a coordinated development effort for applications software, systems software, hardware (i.e., massively parallel machines) and peripherals. It will require testbed systems, a key tool for developing the required software, that are networked for easy access by other scientists. These testbeds will thus serve as a proof of concept and, when scaled up, can be sold commercially to satisfy this country's computing requirements.

In summary, NASA's computational fluid dynamics program, paced by physical modeling and the computing system, has played a key role in the aerodynamic simulation process and has impacted the aerospace vehicle design process through cooperative programs with industry. Such programs require an effective and efficient computer network as part of the computational facility for electronically linking the participating organizations. An optimal computer network is discipline-dependent with gateways to other similar networks, has a high bandwidth and is interactive. NASnet of NASA's Numerical Aerodynamic Simulation Facility is a successful example of a discipline-dependent network that supports collaboration, possesses sufficient data transfer capacity, interconnects effectively with other networks and ensures privacy of use. Progress in computing can be accelerated through the continued pioneering of hardware, software, algorithms, and peripheral technology and the effective commercialization of that technology. NASA's growing emphasis on computational aerosciences and supercomputing technology, with an objective of developing validated applications software today for the computers of tomorrow, will help to ensure this country's preeminence in aeronautics.

Thank you, Mr. Chairman. That completes my prepared testimony. I would be happy to answer questions.

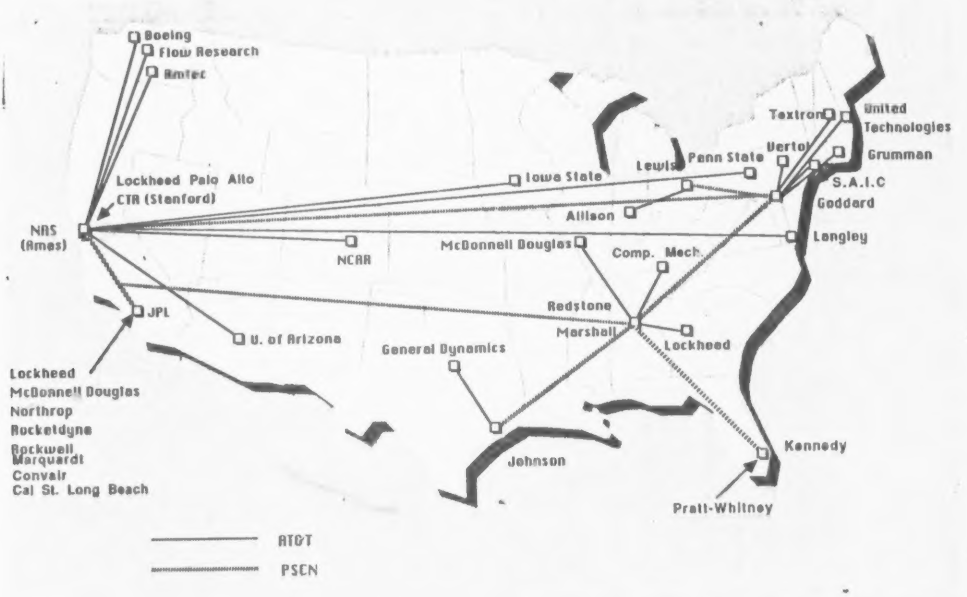
Boeing
Flow Research
Amtec
Lockheed Palo Alto
CTR (Stanford)
NAS
(Ames)
NCAR
JPL
U. of Arizona
Lockheed
McDonnell Douglas
Northrop
Rocketdyne
Rockwell
Marquardt
Convair
Cal St. Long Beach
Iowa State
Lewis
Penn State
Allison
McDonnell Douglas
Comp. Mech.
Redstone
Marshall
Lockheed
General Dynamics
Johnson
Kennedy
Pratt-Whitney
Textron
United
Technologies
Vertol
Grumman
S.A.I.C
Goddard
Langley
AT&T
PSCN

Figure 1. - NASnet User Sites

Senator GORE. Thank you very much.

Our last witness on this panel is Dr. Kenneth M. King, President of EDUCOM in Princeton, New Jersey.

Welcome, please proceed.

Dr. KING. I would like to use this overhead projector if you do not mind.

Senator GORE. Fine.

STATEMENT OF DR. KENNETH M. KING, PRESIDENT, EDUCOM

Dr. KING. EDUCOM is an organization representing 550 colleges and universities in matters concerning information technology and we have 100 corporate associates with us.

EDUCOM was a major sponsor of the National Net '88 in Washington which examined problems associated with national instruction and research networks and we have published a proceeding in this book. Participants of that conference were people from the highest level of industry, government and universities, and these people were unanimous in their recognition of a need for a national network.

EDUCOM also manages the BITNET Information Center. BITNET is one of the current low-speed important national networks, and BITNET connects 400 U.S. colleges, universities and industrial laboratories. It does electronic transfer of mail in small text files. We heard earlier that it is not adequate to transmit big files.

It is the highway equivalent of an electronic trailway. It nonetheless has had significant impact on scientific collaboration. When the supernova explosion occurred over Chile, the scientists around the world had a small amount of time to get instrumentation in place and BITNET and its European equivalent, EARN, played a major role in facilitating that coordination.

It took 7 years of BITNET growth to get one site in every State and, incidentally, I would hope that as we build the national highway equivalent of a network, we can contract that. It became financially self-sustaining in five years and it is growing at the rate of 100 institutions per year.

The universities I represent regard building a national research and instruction network and the knowledge management system on that network as the highest priority. So we believe that a high-speed national instruction research network will enhance productivity by reducing the access time to information, enable scholarly collaboration wherever researchers are located, accelerate the transfer of research results to the private sector and substantially benefit national competitiveness.

What these networks enable people to do is organize all the information on one subject in one place and to give people access to it which significantly decreases the response time to information.

Examples of this are all the information on what to do in case of a toxic spill in one place or all the information on what to do if your daughter drinks some cleaning fluids in one place and there are files that are being created on land management use, on molecular structures, on a whole set of subjects in one place that people can access through the network.

Scholars at our smaller institutions tell me that the existence of even the low-speed network, BITNET has brought them back into the mainstream of their discipline because they can collaborate with researchers at bigger institutions around the world.

The networks themselves illustrate the fact that it is possible to transfer research results to the private sector and this is illustrated by the fact that IBM and MCI are cooperating in building the national backbone because they want to learn how to build and manage a big network.

NYNEX is participating in building and running one of the regional networks, NYSERNET. Bell of Pennsylvania is participating in building the Pennsylvania network. And so we do have industrial cooperation in building this network and they are interested because they are interested in learning how to do it.

Now, I would like to talk to the major problems as I see them. We see the Federal role as providing leadership and planning and stable funding and providing mechanisms for cooperation between higher education, industry and government.

And I might incidentally add that in the last year I have been to Japan and to Europe and those conditions of government leadership and stable funding exist in those countries.

This country has a FRICC, a Federal Research Interagency Coordinating Committee, and I absolutely believe and the people I represent believe that we need a single focal point in this country to build this national highway system. I do not believe the national highway system was built by a committee. And I do not believe we can build the national network with a committee.

So we need to develop a mechanism for providing leadership in stable funding and we also have to assure broad access to the network to scholars and colleges and universities and industrial and government laboratories in every State. Without this participation, 200 universities will connect to the network but we will not get the smaller colleges and universities connected.

In the absence of Federal leadership and stable funding, there are many organizations, and BITNET is one of them, continuing to build these electronic trailways because they do not have faith that the national highway equivalent of the national highway system will be built.

Now, we are talking, I think, about relatively small money in the first phase. The interstate backbone piece of the network, while the backbone itself costs about $15 million a year to build, of which I believe the NSF is providing $3 million and the State of Michigan, IBM and MCI are providing $12 million, there is a need to expand the capacity of that network into the second and third phases through research and development and that has been estimated at $50 million a year. And I might add there is no source of that funding that has been identified.

Connecting a university to the interstate highway system costs 50 to $100,000 per university and if you want to connect 500, that is per year, it costs about $25 million per year.

Now, States, universities and some funding from the government has connected 200 universities to this network but there is no funding source or no incentive to get the next 200 connected.

To connect researchers on a campus to the interstate network system, costs about $500 per year per connection, so if you connect the scholars at 500 universities, the estimate is that that will cost about $250 million a year.

Now, the reason industry is interested in joining us is that there is a huge investment at these universities annually in their equipment to connect to this network and so it represents an extraordinary market for them.

The keynote speaker at National Net '88 said "We believe that national research and education networks, based on cooperation and interaction between government, industry and higher education, are vital for the future. First and foremost, a National Research Network is strategically important to our country's economic future."

She also went on to say "We have a national highway system, a national postal system, a national broadcasting network and a national anthem. It is time for our networks to start singing the same tune." And that requires Federal leadership.

So in summary, universities I represent believe that the nation that provides an effective, high speed instruction and research network first will enjoy a significant advantage. We urgently need a Federal commitment to provide leadership and funding. The most important element is leadership. Thank you.

[The statement and questions and answers follow:]

Statement of Kenneth M. King, President, EDUCOM

I am pleased to address the subcommittee on the subject of computer networking, an area that EDUCOM has been working on for almost 25 years. As President of EDUCOM, I represent a consortium of more than 550 universities engaged in cooperative and collaborative programs involving information technology. In addition EDUCOM has more than 100 corporate associates with a stake in information technology in higher education that help support our activities. The higher education community has defined two major priorities for cooperative activity in the arena of information technology over the next decade. One is creation of a national instruction and research network and a knowledge management system on that network. The other is to develop and infuse into the curriculum information technology programs to improve instruction and learning.

I am here today to address the first of these priority areas. Our goal is to connect every scholar in this country to every other scholar, to national research resources, and to knowledge databases. This goal includes researchers in industrial and governmental laboratories. Ultimately we strive to connect every scholar in the world to every other scholar in the world, so a component of our effort is international in scope.

Networking is making a major impact on the lives and productivity of scholars. First, it has reduced the isolation of scholars at small and remote institutions by enabling them to collaborate with colleagues and to have access to specialized research resources such as supercomputers, accelerators and telescopes, as well as information resources such as national bibliographic databases. Second, it has assisted in bringing these scholars more fully into research activities in their disciplines, which will not only enhance the rate of research progress in many fields but also help to keep instructional programs at small institutions current and help them recruit able faculty. An important role of government support for networking is to ensure the broadest possible connectivity and access to scholars at small and remote institutions. Without government support these institutions will be left out.

Academic networks are essential to the conduct of big science. For example, progress in understanding the fundamental nature of matter involves collaborations among hundreds of scientists at many institutions. Major experiments on accelerators require several years of intensive work by geographically distributed specialists who must coordinate their activities. This coordination is enabled by effective electronic communication. When a rare event in the heavens occurred last year, a supernova seen over Chile, an international effort to get experiments in place around the world in a matter of hours was greatly facilitated by the BITNET network in the United States and Asia, and its European counterpart, EARN.

Studies have shown that productivity is proportional to access time to information. By greatly reducing access time, networking is a critical element of the national competitiveness and productivity equation. The explosive growth in overnight mail services, fax services, and corporate and university electronic networks over the past five years is testimony to the importance of reducing access time to information. Our vision of a national network would enable text, digitized sound, and pictures to be transmitted from one place to another in seconds rather than days. Scholars at colleges and universities, industrial laboratories, and government laboratories would be able to work simultaneously on a common body of knowledge. Thus the network will accelerate the transfer of research results to the private sector. The impact of this technology on scholarly productivity and national competitiveness will be profound.

Networks have made it possible to assemble all the current information on a research topic in one place and to make this information available electronically. Examples include databases dealing with toxic spills and toxic substances, land use management, chemical and molecular structures, as well as emerging databases on AIDS and high temperature superconductivity research results. Being able to access almost everything that is known on a subject electronically and developing a consistent scheme for assembling information from many sources on the network in order to develop new information are major goals of the national networking effort. The impact of this capability on all sectors of society will be extraordinary. Thus all of us have a major stake in the development of this technology.

The importance of this technology coupled with the absence of an effective coordination mechanism has resulted in a rapid proliferation of networks that do not effectively communicate with each other, but that nevertheless have had an enormous impact on scholarship. EDUCOM manages the network information center for one of these networks, BITNET (Because It's Time network). This network, started in 1981 by the City University of New York and Yale University, has grown to connect more than 400 colleges and universities in the United States and more than 800 colleges and universities worldwide. To join BITNET, an institution must lease a communications line to an institution already on the network and agree to provide at least one connection for another university. Thus the network has grown and is continuing to grow at the rate of more than 100 institutions per year worldwide. This network costs each institution about $25,000 per year to sustain. The success of BITNET was greatly facilitated by a grant from IBM to EDUCOM and the City University of New York. A similar IBM grant helped develop BITNET's Canadian, European, and Asian counterparts. These networks are now self-sustaining and connect thousands of computers from virtually every manufacturer.

The low speed BITNET network supports the transmission of electronic mail and small text files. It does not support interactive connections to knowledge databases, nor does it enable the transmission of digitized sound (voice, music, etc.) or pictures. An additional problem is that the BITNET technical protocols are incompatible with most other academic networks, which requires users to negotiate complicated and often unreliable "gateways" between networks.

The current situation is not unlike the situation that might have developed if the states independently had responsibility for planning and developing their highway systems so that major highways crossing one state might not connect to highways in adjoining states except through side streets. Signage might be different from state to state. Interstate traffic and commerce would be stymied.

In an effort to create coordination at least in higher education, forty institutions of higher education have contributed funding and people to an EDUCOM Networking and Telecommunications Task Force, which has been working on all of the problems of creating a national network. This group quickly realized that our dream of a true, high speed, national network for higher education (a digital superhighway) could only be realized through cooperation among higher education, government, and industry.

In April of this year, EDUCOM, the National Association of State Universities and Land Grant Colleges, and the New York State Educational and Research Network sponsored a major national networking conference, held here in Washington with the assistance of the National Science Foundation. The proceedings of this conference, contained in the current issue of the *EDUCOM Bulletin*, provide the most up-to-date statement of government, industry, and educational perspectives on the national network. Copies of these proceedings have been made available to the subcommittee members.

Speakers at this conference included representatives from industry, such as Ellen Hancock, Vice President and General Manager, Communications Systems, at IBM; A.G. Fraser, Executive Director of the Research Information Sciences Division at AT&T Bell Laboratories; and Robert E. Kahn, President of the Corporation for National Research Initiatives. Representatives from government included Congressman Douglas Walgren; Erich Bloch, Director of the National Science Foundation; and Henriette D. Avram, Assistant Librarian for Processing Services at the Library of Congress. In addition there were many key people in higher education, networking specialists as well as senior administrators, such as Dr. Donald Langenberg, Chancellor of the University of Illinois-Chicago and Dr. Harold Raveche´, President of Stevens Institute of Technology.

These representatives shared the view expressed by Ellen Hancock in her keynote address: "We believe that national research and education networks, based on the cooperation and interactions between government, industry and higher education, are vital for the future. First and foremost, a national research network is strategically important to our country's economic future." She went on to say, "We have a national transportation system, a national postal system, national broadcast networks, even a national anthem. It is time to get our scientific networks singing the same tune, too."

The theme of cooperation among government, industry and education and the challenge we jointly face is depicted by the picture on the cover of the *EDUCOM Bulletin*. It shows three men trying to ride a unicycle down a communication line to a stable platform. The man driving the unicycle is from government. On his shoulders, and thus able to see farther, is the representative from industry. Riding his shoulders, and thus able to see farthest of all, is the representative from higher education. Success depends on coordination and cooperation of these elements. The key to stability and the driving force is government leadership. Unfortunately, the current state of affairs is that government support and leadership are unstable, with the possibility that the whole cooperative effort may fall apart. I would like to focus now on the critical support and leadership role of the government in creating a national network that will enhance scholarly and industrial productivity.

The National Science Foundation has taken one major step toward a national research network by its funding of a contract with MERIT, Inc., for the NSFNET. This high capacity backbone service network is managed by a joint venture of IBM, MCI, and MERIT (which is an academic network consortium in the State of Michigan). This backbone connects to mid-level or regional networks at 13 places and costs about $15 million per year. NSF provides $3 million per year and IBM, MCI, and the Michigan Strategic Fund jointly provide about $12 million per year for the project. This constitutes a major example of government, industry, and university cooperation and of government's catalytic potential for leverage. The backbone, a successor to a smaller, lower speed backbone network, became operational at 1.5 million bits per second on July 1, 1988.

Mid-level networks connecting more than 200 colleges and universities have formed and are being formed. These high-speed networks have been supported in part from the National Science Foundation and in part by state, corporate, and university contributions. It costs of the order of 50 to 100 thousand dollars per year per institution to support communications costs to connect to a mid-level network, not counting campus networking and support costs. As yet there is no stable, continuing source of funding for these mid-level networks and, beyond the 200 best-funded colleges and universities, no identifiable source of funds to connect the next 200 colleges and the next beyond that. Together, the backbone and mid-level networks constitute the interstate network system.

Building a campus network connecting scholars to the mid-level networks costs about $500 per year per faculty or student network user. Many campuses are building these networks largely with their own funds, because they have faith that funding for the interstate networks is stable and will be available to upgrade them as demand increases. They must have faith also that the policies and programs permitting the development of important services on this network will fall into place. The aggregate annual cost of connecting the scholars at 500 universities to the emerging network is estimated to be $250 million per year. The EDUCOM Network and Telecommunications Task Force has just published a case study book on successful campus networking projects at 10 campuses. Copies of *Campus Networking Strategies* have been provided to the members of the committee.

In order to upgrade the interstate networks as institutions develop their local networks, research and development in networking is required. No source of funding for this effort has been identified. The 1987 Office of Science and Technology Policy (OSTP) report, *A Research and Development Strategy for High-Performance Computing*, estimates the annual cost of this R&D for communications at $50 million per year. In time users of these networks can be expected to support their cost, but this will certainly not be the case during the next phase of development. In fact when the network is operational and reasonably mature, the users should pay its costs. We estimate that this will require five to ten years; BITNET took five years to be self-sustaining.

Thus, federal government funding requirements may be of the order of $100 million per year for five to ten years and should serve to build the interstate portion of the national education and research network. A major component of government support is required to ensure broad access and the stability of the network during its developing years. It took BITNET seven years to get at least one institution in every state connected. I would hope that government support would ensure this level of connectivity for a higher-speed network in the next two years. A major current problem is the lack of stability of funding and a plan to insure that if campuses put their funds into a local area network, a connection to the national network will be available when they need it.

Aside from funding, a critical requirement is government leadership. Five federal agencies—NSF, DOD, NASA, DOE and HHS—are involved in supporting networking and are members of a federal research interagency coordinating committee (FRICC). I do not believe that a committee can provide leadership, although it can inform and support leadership. I believe it is critical that the Congress and administration define an effective leadership function. One possible mechanism would be to ask OSTP to define this leadership function. OSTP could ask the FRICC to define a leadership mechanism and protocols defining the role of the key agencies in this effort. Without leadership and planning and stable funding, the whole effort could collapse. It is urgent that the leadership issue be resolved immediately.

There are many other important issues, but the key ones are leadership and funding during the development phase of this effort. Leadership from the federal government is as essential to this effort as it was to developing the federal highway system. In the past year, I have traveled to Asia and Europe to talk to university people building networking support. The stability of funding and the leadership of government in this area are not an issue in Japan and in many countries in Europe -- Germany, France and England in particular. I believe that the first country to provide an effective research and education network will have a significant advantage. To achieve this advantage, we urgently need a commitment from the federal government.

Answers to Questions from Senator Ernest F. Hollings
Chairman, Committee on Commerce, Science and Transportation

Kenneth M. King
President, EDUCOM

1. The BITNET network now serves tens of thousands of computer users across the country.

> *Who are these people? Are they mostly scientists? Do students, as well as faculty, have access to the network?*

The network does not keep data on who uses the network. Accounts providing access are distributed by the member institutions of Bitnet. Bitnet currently has 400 member institutions in the United States with at least one institution in every state and this network is connected to about 400 institutions world-wide (who are members of cooperating networks which provide the same services and have compatible policies) in Canada, Europe, Asia, Africa and South America. The rules of Bitnet permit an institution to provide access to faculty, administrators, and students for uses consistent with the academic purposes of the institution. Some institutions do not permit student use. If one examines telephone directories and business cards where Bitnet addresses have begun to appear, and talks to Bitnet institutional representatives on campuses, one concludes that the faculty constitute the largest segment of users because they are involved in the largest number of collaborations and they come from all disciplines not just the sciences. However, faculty from the Sciences and Engineering constitute the most active group because they have a long history of collaborations and they are the most computer literate segment of the faculty. A large majority of the 1,557 computers connected to Bitnet in the United States (many members have multiple computers connected) are operated by faculty in colleges and departments in support of scientific research.

> *Where are these people located? Are they mostly at the big research universities?*

The available evidence indicates a pattern where Bitnet usage at an institution depends on the level of computer activity on that campus and the length of time the institution has been connected. The big research institutions connected earliest (Bitnet is seven years old) and usually have a high density of computers on the campus so they have a high level of activity. However, computer intensive liberal arts colleges and comprehensive state and private institutions are also very active. In discussions with faculty, one gets the impression that networking has changed the quality of life at remote and isolated institutions. For example, President Jeffrey Holland at Brigham Young University said: "Networking is very important to higher education but it is absolutely crucial to the isolated universities. Networking removes them from isolation, gives them access to excellent library and computing resources, and provides opportunities for scholarly interchange which makes it possible for them to compete on a much more equal footing with the best universities for faculty and students."

2. *Most users of BITNET are at universities and colleges. Yet, it's clear that researchers in industry can benefit from national computer networks as well.*

Are there any industrial users on BITNET?

The rules of Bitnet permit industrial researchers to connect to Bitnet for the purpose of pursuing collaborative activities with higher education consistent with the academic purposes of institutions of higher education. All commercial traffic is prohibited. Members of Bitnet include in addition to colleges and universities, government research laboratories; not for profit research organizations and foundations with connections to higher education; and about 10 industrial research laboratories which sponsor research on campuses. In addition, there are gateways to Bitnet from CSnet; the Internet; and a range of other networks which accept industrial organizations as members. Bitnet accepts traffic through these gateways which is consistent with the academic purposes of Bitnet.

Do you have any plans to increase the number of industrial users?

Bitnet has not solicited industrial connections. Bitnet is currently pursuing merger with CSnet (Computer Science network). This merger would bring about 40 additional industrial organizations into ONEnet, the merged organization. CSnet connects computer science departments in higher education with computer scientists in industrial corporations pursuing collaborative research.

What are the pros and cons of having university and industry researchers on one national research network?

There are a rapidly growing set of relationships between researchers in universities and industry that have been encouraged by network connections. These relationships are very valuable to colleges and universities and apparently to industrial organizations as well. The only problems that have emerged are rare appearances of commercial traffic on the academic network or more accurately traffic in a grey area that might be interpreted as commercial and occasional problems with computer viruses introduced into industrial networks through the gateway from academic networks. The overwhelming consensus appears to be that the value of these connections greatly exceeds the occasional problems.

3. There have been a number of articles recently on computer viruses and the dangers they pose. Because computer viruses can spread so easily through a computer network, they are a particular threat to computer network users.

Have computer viruses been a problem on BITNET?

Yes, they have been and continue to be a problem. There was for example a major disruption of Bitnet caused by a Christmas greeting virus which appeared on December 18, 1987. The virus appeared in the form of a mail message which requested recipients of that message to execute a program module included with the message to receive a Christmas greeting. When executed, the virus replicated many times, and the resulting traffic clogged the network for about two hours. No files were destroyed. This virus educated network users to the rule that you never execute an unsolicited program regardless of source.

Do you think they will pose serious problems to researchers on the National Research Network?

Viruses will continue to be an occasional problem on the network, but, I expect the level of the problem will decrease over time as the computers connected to the network develop counter virus measures equivalent to the immunity response of humans to viruses. In the early days of shared computers on campuses, colleges had problems with student hackers who exploited holes in commercial operating systems. Vendors then took responsibility for closing those holes and the magnitude of the problem has decreased dramatically. The November 1988 virus on the Internet exploited holes in a non-commercial operating system widely used by the research community. With coordinated and serious attention to closing those holes, I expect the level of the problem on the Internet will decrease.

Do you think that there will be many researchers or businesses that will be reluctant to hook up to a National Research Network because of their concern about computer viruses?

This could be a serious problem if the Internet community does not step up to the problem and develop an arsenal of immune responses to externally distributed viruses. However, these networks have become essential to colleges and universities and industrial networks have become essential to a growing number of corporations. The value of the communication links connecting them greatly exceed the risks in the view of those who depend on these links.

4. NSFNET and the gigabit network proposed in the OSTP report are to give researchers access to supercomputers and carry electronic mail. I understand that designing such a dual-use network is not too difficult. However, there are some inherent conflicts. Researchers need an electronic mail network that reaches as many people as possible. On the other hand, we may want to restrict who has access to supercomputers.

Just how difficult will it be to have one network which provides both supercomputer access and electronic mail service?

From the perspective of faculty, it is highly desirable to have network access to specialized computing resources such as supercomputers; specialized instruments such as particle accelerators and telescopes; and information data banks such as libraries; in addition to mail and message connections with their peers and collaborators. Since all of the networks are being connected by gateways to enable these possibilities, the only feasible mechanism for restricting access to a particular resource on a network is at the resource itself. The supercomputer centers have in my view established reasonable procedures to restrict access to authorized users. These procedures need to be audited and monitored to guarantee their continued effectiveness.

Is BITNET connected to the supercomputer centers now? Why or why not?

Bitnet is connected to the supercomputer centers and enables researchers not connected to the Internet to pursue research at the supercomputer centers. Researchers at these institutions use dial up connections to interact with the supercomputers in real time if necessary but submit programs and receive output from the supercomputers via Bitnet. Bitnet does not support logon to computing resources on the network, as does the Internet, due to limited bandwidth. However there are about twice as many institutions connected to Bitnet than there are to the Internet. Many of these are small and remote and Bitnet provides an essential connection.

5. In your written testimony you discuss the problems of the regional networks that link individual campuses to the NSFNET backbone. Because of funding problems, these networks are often the weak link in the network.

Why are the regional networks having problems securing funding?

The cost of a regional connection ranges from $50,000 to $100,000 per year per institution depending on how remote an institution is from a connecting point. In addition, an institution has the problem of building a local area network connecting its faculty to the point of connection to the regional network at a cost of about $500 per year per connection. These costs are prohibitive for small and more remote institutions and particularly those with no significant cadre of faculty demanding supercomputer access. Since the faculty at these institutions have no access to supercomputers and other national research resources, they haven't used them and because they haven't used them they haven't demanded access in numbers that create a priority for scarce institutional resources. Regional networks whose boundaries are confined to a single state such as Nysernet, in New York State and PREPnet in Pennsylvania have been successful in getting State dollars to partially match Federal dollars so that many colleges and universities in these states are connected and these networks have also enjoyed some industrial cost sharing. Regional networks like Midnet and Suranet, which span a number of states, connect only the major research universities in those states and do not enjoy state funding to a significant degree. State and Federal funding are essential to insure broad connectivity of small institutions in these states. It has been historically true that new services require subsidy during their early years until the benefits of these services are apparent. It is my belief that the Federal government should take responsibility for providing at least one connection in every state and provide incentives to states to support broad connectivity within the state during the formative period of the network. This period is likely to last five to ten years.

> *I would think that state governments would recognize the economic benefits that high-speed networks can provide. Why aren't they chipping in more?*

Major industrial states have clearly recognized these benefits and are chipping in where the networks are organized to focus these benefits on a single state. State science and technology agencies or economic development agencies have not seen an incentive to contribute to multi-state regional networks which include states with which they compete. Federal leadership and some funding could fix this problem.

> *Given that NSF has a limited budget for networking, do you think it was wise to concentrate on the high-volume backbone that connects the supercomputer centers?*

I believe this decision was both wise and essential for the following reason:

The only incentive institutions have to connect to a high speed network is to achieve interactive access to an important shared resource. The only important interactive shared resource currently available to a large number of scientists is the national supercomputers. Thus, it was essential to connect to them first. As libraries and databases get connected to the network, more and more institutions will have an incentive to connect to the network and the number of backbone connections will need to be increased. A network architecture that focuses on connections to important high volume shared resources on the backbone is optimal from a traffic perspective.

6. BITNET is funded by payments from each institution on the network. Other networks are paid for by the federal government. Still other networks are paid for by individual users.

> *What are the advantages and disadvantages of each of these funding schemes?*

This is a complex question. Where benefits clearly accrue to individual users and there are no disadvantages to encouraging multiple providers whose success depends on responsiveness to user requirements, a fee paid by individual users is optimal. For services which are essential to the institutional infrastructure and where equality of access is important, institutional funding is essential. For example, great libraries would not have developed on campuses if users paid for access to libraries and paid a lease fee for each book borrowed. Bitnet charges by institution and not for use by individuals because its members believe that its services provide an essential infrastructure for shared scholarship and equality of access is essential. Over time this situation may change but the Bitnet Trustees believe that for the foreseeable future some organization should take responsibility for developing and maintaining low cost network services that are universally available to higher education and they have assumed that responsibility.

The major arguments for federal leadership and funding of the national high speed research network include:

1. Like the national highway system, a high speed network is an essential component of the national research infrastructure because it provides access to expensive shared resources;

2. the value to the public good of researchers communicating, collaborating and sharing information and improved research productivity and technology transfer between universities and industry and other benefits justify the cost to the public (i.e. the value exceeds the cost);

3. the nation's telecommunications infrastructure is so important to trade, productivity and growth that it is in the national interest to leverage the computer research community to produce pre-proprietary prototypes of network features of potential commercial value; and

4. the needs of the research community exceed and anticipate

> needs from the commercial sector so that market incentives on the commercial side will not address the essential requirements of the research community without some public support at least in the short term.

Long term, the commercial sector could respond to the needs of the research community in a way that makes the advantage of responsiveness to user needs important enough to reduce the requirements of federal, state and institutional support of network infrastructure. It is my view, however, that without federal funding and leadership in the short term, the national research network will fall far short of its potential to the detriment of the public good.

Senator GORE. Thank you very much. That is a very thoughtful statement. I appreciate the insightful remarks of everyone on this panel.

Dr. Kahn, what about this question of user fees? If researchers and other users are not charged for using the network, will there not be an unlimited demand for service and great difficulty in allocating capacity?

Dr. KAHN. I would like to approach the user fee question from the point of view not of just the research community alone but the broader context in society to which this kind of a network might be placed. That is where I think the real self-supporting activity can come from. Otherwise, all we are talking about is a redistribution of Federal funds which go into the Federal research projects.

If you have an adequate support base, you can provide the capacity to make an effective network. This will require user fees, even if the Government provides some of the funds via research contracts.

We need to find a way to start by generating a way for the research community to gain linkages to those parts of the non-research community with which they need to interact. That happens all the time in other domains. We do not have a separate telephone system just for the research community and researchers interact by telephone with people who might not normally be considered researchers.

I think there will be ample——

Senator GORE. And they pay for it?

Dr. KAHN. They pay for it out of Federal budgets, out of overheads, in various ways. And I think that will continue. We can make use of the normal mechanism for funding academic research, that is largely Federal. In addition, we can attract some industry money as well.

As you broaden the research net or allow for expanded activity beyond the research community per se, other organizations will be willing to pay for their participation. It will be viewed by both parties as a valuable interconnection. In the long term I think that can serve as a source of added revenue which will make the whole system self-supporting.

Senator GORE. You mentioned the digital library initiative as one example of infrastructure above the network itself that would be essential to fully utilize the new capacity.

You said you had some other examples. I do not want you to give them all but give us just a couple of others.

Dr. KAHN. Okay. Another one that I think would be useful to start within the research community, but with obvious implications to the well-being of the country, is an electronic transaction framework. These are the mechanisms that allow for the computers to interact with each other. These mechanisms will allow people to go into open bidding mode, find out what is available, get back responses and negotiate for deliveries, take deliveries, perhaps electronically, perhaps not electronically, to validate deliveries when they have been taken, and to deal with issues of accounting. There is probably no single action that can be taken to improve productivity more up front than just getting rid of a lot of the paper and personnel overheads associated with procurement.

Such a capability has very broad implications everywhere and I think it is a leverage point into the whole area of custom design and flexible manufacturing. You can envision a future in which a custom design can be done in a workstation—where the implementation plan can be created right there in the same machine and bidding can be done right then electronically on a national network. We can get estimates of delivery time and manufacture from everybody in the country who is potentially capable of dealing with that problem and interested in responding. That is one example.

A second one which is particularly interesting to me is the idea of developing some kind of a medical exchange in which access to medical information could be made broadly available, not only in the research community where people are doing medical research but also eventually outside the research community where individuals would like to have sources of contact with people who have had perhaps similar experiences with medical problems. Or perhaps just to communicate with other people who may be able to assist them. These could be health vs. medical interactions that are social as opposed to highly professional, but I think it could support both.

Senator GORE. Very good.

Dr. Dickinson, I understand that the work underway to model the global climate is generally recognized as requiring larger data sets and more computing power than almost any other scientific endeavor. Is that correct?

Dr. DICKINSON. Well, I cannot make comparisons with other areas but I certainly know that we have very large data sets which, in thinking in terms of computer tapes, that we do simulations that end up putting the output on equivalent to hundreds of computer tapes and after one does many simulations, it is easy to stack up over a few years data equivalent to many, many thousands of computer tapes.

Senator GORE. Are the existing networks important to your collaborative research?

Dr. DICKINSON. We are anticipating they will be in the future.

Senator GORE. They are not now?

Dr. DICKINSON. They have not been crucial up to now. They have been important from the viewpoint primarily of communication rather than exchange of computer output. For my own research, I have many colleagues at universities that use our computers that come in on networks and have been doing this for close to a decade now, and to them it is crucial, though before they did that they had alternatives. They came and visited us for the summertime and did their work on-site, which was very inconvenient for them but they still could do their work.

I see networks as something that is a very important convenience but one that should be economical to use. Scientists spending a budget should be able to choose cheaper alternatives, if available, rather than networks under some circumstances.

Federal Express was mentioned earlier. For many data exchange activities, scientists might be happy to get overnight service. They could use some such service as Federal Express as an alternative so that networks, to be competitive for those satisfied with overnight service would have to have a pricing that was not huge compared to the cost of Federal Express.

Senator **Gore**. Well, as you may know, I chaired the first hearing on global warming in the Congress some eight years ago and since that time I have been pushing very hard for action to address the problem. And I am convinced that this proposed network will be important to NCAR's work in this area.

Dr. Kutler, what percentage of the computer time on the Numerical Aerodynamics Simulator is used by NASA researchers at Ames? Do you know offhand?

Dr. **Kutler**. I think it is approximately 23 percent of the computer time is used by the NASA Ames scientists.

Senator **Gore**. Approximately 23 percent?

Dr. **Kutler**. Twenty-three percent of the time.

Senator **Gore**. Can't you be more precise than that?

What percentage of users access your supercomputers over NASnet?

Dr. **Kutler**. All of the out-of-house users, which would be the balance of users that are not using locally, which would be NASA Ames, so the balance is 77 percent, I would imagine, of the users use the NASnet network.

Senator **Gore**. We certainly want to congratulate NASA on its successful test of the Space Shuttle yesterday. I know the whole agency is greatly relieved and excited about that.

You mentioned replacing wind tunnel tests. That generates huge amounts of data. Simulating a wind tunnel test is extremely complex, is it not? I think you said it requires 25 hours on a supercomputer?

Dr. **Kutler**. It took 25 hours to produce that calculation using the supercomputer.

Senator **Gore**. What is the significance of this three-dimensional image? It is neat but tell me what it is.

Dr. **Kutler**. Well, we use the three-dimensional viewer to help provide the scientists who have to analyze the data that they get, to do it in a very quick fashion, very efficient fashion.

We actually had to look at the computer output which is a very inefficient way of trying to understand the solution of what is going on. This is just a means of getting that understanding very, very fast so that we can make changes to the code and make changes to the flow conditions to determine exactly what happens in these flow fields.

Senator **Gore**. Dr. King, I enjoyed your testimony. What gets sent over BITNET now besides electronic mail?

Dr. **King**. You can just send small text files and messages. The messages are interactive and if two scientists on each end of the network are at their terminals, they can send messages electronically. One of the common messages is, "Are you free? ", which indicates that if you are, I would like to get you on the telephone. So it avoids telephone tag.

The other common sort of message is, "Does anyone know, has anyone encountered the following problem? And do you know how to deal with it?" Other things, newsletters on emerging technologies are in a database on the network and people access that database to find out what is going on. There are a lot of those.

Senator GORE. But they do not send graphics or data?

Dr. KING. No. It does not have the capacity to send digitized voice or images. That is why people use fax and overnight mail.

If you had a high-speed national network, you could transmit those things over the network and you could get access to information in seconds rather than 24 hours or 36 hours or days.

Senator GORE. Now, you referred to the work in other countries to complete networks of the sort we are discussing here. Could you give us some examples of what is being done overseas and examples of where the United States is falling behind?

Dr. KING. Well, networks are being built in Japan, France, Germany, England, and in other places. When you talk with people who are working on those networks, they complain about inadequate funding but they do not complain about government leadership and the stability of the funds that they have.

So they are busy building networks that will replicate our 1.5-megabit backbone network. They have in place lower-speed networks that are much more universal. They have been developing management techniques and it is the government/industry/university combined effort which is coordinated that is moving them ahead of us in my belief.

Senator GORE. Which country is farthest ahead, Japan?

Dr. KING. In the high-speed networks and in fax transmission, they suffer a handicap in the low-speed networks. They are of no use to them because Kanji is a difficult thing to enter through a keyboard. So they are very dependent.

Senator GORE. So far.

Dr. KING. That's right. So they are very dependent on image transmission systems, fax systems, and so they are working on high-speed systems because they are dependent on them.

Senator GORE. They are farther ahead?

Dr. KING. Well, it is hard to determine. They are not complaining the way the people I represent complain.

Senator GORE. I have a number of other questions in writing that I hope you all will be willing to respond to. We are running short on time so we are going to have to move on but I really enjoyed this panel and I do very much appreciate your participation here today. We are going to pursue this vigorously. Thank you.

Let me call our last panel today to talk about building a National Research Network. Dr. John Connolly, Director of the Center for Computational Sciences and Professor of Physics at the University of Kentucky in Lexington; Dr. William Wulf, Assistant Director of the Directorate for Computer and Information Science and Engineering at the NSF; and Dr. Alan McAdams, Associate Professor at the Johnson Graduate School of Management at Cornell.

I welcome all of you. Thank you for joining us today. I apologize for the length of the morning's session and thank you for your patience.

We will start with you, Dr. Connolly.

STATEMENT OF JOHN CONNOLLY, DIRECTOR, CENTER FOR COMPUTATIONAL SCIENCES, AND PROFESSOR OF PHYSICS, UNIVERSITY OF KENTUCKY

Dr. CONNOLLY. Thank you very much. I am very happy to be here and in the interest of time, I will be very brief. I just have a couple of points to make and a recommendation and I brought a short movie with me which I think you will enjoy which has to do with an actual research application on an actual network using graphics, graphical visualization of supercomputer results. I think it is important to think that as we are building our electronic interstate, we want to be careful to avoid a massive graphic jam.

Senator GORE. I had another witness like you one time.

Go ahead.

Dr. CONNOLLY. I bring kind of a different perspective from the other panelists because I was there at the beginning at NSF when we set up the supercomputer centers and the network, so I know a little of the history of where we have been, where we thought we would be going and where we are now.

I have three points to make. The first is that the NSF programs in advanced scientific computing have been very successful, beyond what our original expectations were.

The second point is that the NSF programs are now in jeopardy because of the limited amounts of appropriation that have been assigned to those programs in the last 2 years.

The third point I want to make is that other solutions to the funding dilemma, such as charge-back mechanisms, relying on State and local funds, soliciting industrial support or expecting the mission-oriented agencies to pick up the difference could change the nature of the NSF programs in unsatisfactory ways.

Let me expand on the first point, namely, those of us who are familiar with the scientific literature have seen a recent explosion of articles in almost all branches of science, which would not have been possible if the supercomputer program had not been there. Some of these have filtered down into the popular media and I would not go into detail but there are many applications, a lot of them are important, all of them are interesting. That could be the subject of a whole other hearing.

However, we found the centers have been a victim of their own success. As they have become more and more useful and more popular, they have become saturated. At Kentucky when we were connected into SURANET we were delighted that it gave us access to supercomputer centers all over the country.

Soon, though, we found it was almost impossible during the daytime to get on the network because there were so many other users. Now, that situation has improved somewhat as of July 1 as NSFNET has been upgraded but we know that is going to happen again.

Perhaps it is interesting to go into a little history of the NSF program. We set it up in three phases. The first phase started in 1984 where we were using existing resources. We went out and bought time from commercial vendors. We used some of the independent centers like in Minnesota, Purdue, and Colorado State. We used existing networks like ARPANET and MFENET and that was an essential part of getting the program off the ground.

The second phase was that we created an entirely new national computing resource consisting of the national NSF centers and NSFNET. Originally we had a fairly ambitious program. We had planned on 10 to 12 centers. We had to compromise on five.

The networking was slower than we thought it should be but at the present time, after a remarkably short time as Washington programs go, these centers have established themselves in leaders in computational research in showcases of new technology; NSFNET has provided connections for more than 200 universities and laboratories.

One of the principles we used in designing it was to use state-of-the-art technology but not beyond.

The idea was to have a system that would be functional in a minimal amount of time. If you go too far beyond the state of the art, it takes years before you get something that really works.

A second principle we used in the design of the system was top-down management, namely, that budget was taken off the top of the NSF appropriation rather than accumulated from the bottom. We found that it would just be impossible for 100 different NSF programs to agree on a single unified approach.

I think there is a lesson there if we are trying to make an overall government network.

The NSF program has made a lot of contributions to education and training. It has also had an unexpected result—it led to proliferation of independent supercomputer centers around the country and of statewide networks. NSF provided seed money that showed people the way to go.

That is what leadership is all about. My own center at the University of Kentucky would not have occurred if it had not been for the fact that we had faculty/students on the campus that pressed the administration through their experiences with the NSF centers.

Originally we also had a Phase III program which was part of the plan to advance the state of the art in supercomputer technologies. Part of that was upgrading the centers to the next generation supercomputer. Part of that was an upgrade into the next generation of supercomputer networks.

The Phase III program really never got off the ground. Again, the budget limitations forced us to concentrate on what we had, rather than going ahead.

So that is essentially my second point, that inadequate budgets, if we are forced to stay at the current level, will inevitably lead to obsolescence and the NSF will lose the leadership which it has exerted and lose the momentum which they painstakingly built up over the last four years.

I recommend, of course, the OSTP report as mapping out a plan for building on the past success of these programs and for revitalizing the national program.

I would like to make one last point, I think, that we have to be very careful how we design this. The ideal thing would be to have a lead agency which has power over the planning and the budget process. That may not be the most appropriate mechanism, but if that is not possible, then my second choice would be to have a very strong central planning process, perhaps at the highest levels at the Executive Branch in OSTP.

My own personal opinion is that it should have a broader and more responsible charter than the current FCCSET committees. But the centralized planning is absolutely essential. Otherwise you are going to fractionate into a number of incompatible programs.

The only problem about building the network is achieving the increased capacity and the technology is really very close to achieving the rates for future applications. And all that seems to be required is the will to do it.

And the success of the venture is virtually guaranteed based on our extrapolations from past experience and it is also generally accepted as such a supernetwork will be the major communication medium for the knowledge based for our information society which has been commented on today.

So I will end my testimony there and if you could turn on the movie, not quite right away, I should explain a little bit of what it is about.

This is an actual research project on a very timely topic, namely simulation of water resource management. It is based on research by a professor of engineering in Kansas University who uses the Cray supercomputer at the National Center for Supercomputer Applications in Illinois. I want to thank the Illinois Center for having prepared this movie.

What you will see is an actual research project with the kind of scientific results you can get at four different level network speeds. The first level you will see is the bad old days when all you had was 2,400 bits per second over a telephone line. The second level you will see is the present system, namely 56,000 bits per second over NSFNET.

The third level you will see is what is possible at T-1 speeds which is the first step of the proposed upgrade, and the last example you will see is the kind of science that this professor can do at T-3 speeds which is the second stage, T-3 means 44 million bits per second. You can turn on the movie.

[A movie was shown.]

[The statement and questions and answers:]

STATEMENT OF JOHN CONNOLLY, DIRECTOR, CENTER FOR COMPUTATIONAL SCIENCES, UNIVERSITY OF KENTUCKY

Senators and Fellow Panelists:

In giving this testimony today, I would first like to commend this committee for their past support of federal funding of computer-related research. Without the assistance of this and other Congressional committees, much of the recent progress achieved by NSF and other federal agencies would not have been possible.

Today, I want to share with you some of my views of past progress and future opportunities in the computational field, both from the perspective of my current position, as head of an "independent" supercomputer center, and of my previous position as Director of the Division of Advanced Scientific Computing at NSF. This was the part of NSF which was responsible for initiating and managing the supercomputer center and national network programs.

Briefly, I would like to make three points relating to the NSF and other federal programs, and then make some recommendations for the future, especially with respect to establishing a research and development strategy for high performance computing (the National Computing Initiative):

(1) The first point is that the NSF programs in advanced scientific computing have been very successful, even beyond the original expectations. At the start of the supercomputer initiative, there was a threefold goal to promote supercomputing in research, education and industry. Significant progress has been made in all three aspects, as I will elaborate further later in this testimony.

(2) The second point is that the NSF program is now in jeopardy because of the limited amounts which have been appropriated recently in the NSF budget. The small increments over the past two fiscal years have not allowed the supercomputer centers and network programs to honor their past commitments, let alone advance the state of the art technology. These original commitments were made at a level to provide a minimum amount of service. However, as the technology of supercomputers and networks has developed, many opportunities have arisen, which have been missed, and which will continue to be missed if these programs are not allowed to expand, rather than to contract. Unfortunately, in times of tight budgets, funding agencies like NSF become conservative in their funding strategies, opting for maintaining existing structure at the expense of new initiatives.

(3) The third point is that other solutions to the funding dilemma, such as charge-back mechanisms, relying on state and local funds, soliciting industrial support, or expecting the mission-oriented agencies to pick up the difference, would change the nature of the NSF programs in unsatisfactory ways.

Let me expand on the first point,- namely, the success of the NSF program. Those of us who are familiar with the scientific literature have seen a recent explosion of articles, in almost all branches of science, which would not have been possible without the access to supercomputers provided by the NSF. Some of the results have appeared in the popular media, and you may already have seen some of them. They are many, important and interesting, but there is no time to go into detail here.

The NSF clearly established a leadership role for itself in computers and networks by the establishment of the Advanced Scientific Computing (ASC) initiative in 1984. This office started the NSF supercomputer centers and the NSFNET. The initiative originally envisaged the networking part of the budget to be primarily for access to supercomputers for the scientific computing community. However, it soon became apparent that there was need for more research in new types of networks, so the Division of Networking and Communications Research and Infrastructure (NCRI) was split off from ASC and augmented with computer science initiatives. This split caused consternation in the scientific computing community, as there was the perception that networking funds were being shifted from their community into the computer science community. This tended to exacerbate the tense situation between the communities.

However, if NSF/NCRI were to delineate clearly the funding strategies in the two fields, as I suggested above, then this could go a long way towards reducing the counterproductive competition between them.

Q2: The OSTP report lays out a general timetable and budget for building a 3 gigabit national research network.

- Do you think the OSTP timetable is realistic?

- Specifically, how soon do you expect a 3 gigabit network to be up and running and how much will it cost?

A: The estimates in the OSTP report for attaining the ultra high speed national network are divided into three stages:

(1) The first stage would upgrade the existing research-support networks to T1 speeds (1.5 megabits per second) at a cost of $5M per year for three years. This is certainly possible within existing technology, and is already underway. It could easily be achieved ahead of schedule and within budget.

(2) The second stage would upgrade to T3 (45 megabits per second) for 200-300 research installations, with operating costs of about $50M per year. The second stage would begin to replace the first stage about three to five years after the beginning of the program. Until its installation, this stage would require $5M per year of development costs. Within the uncertainties of technology, this is a reasonable schedule and cost estimate. At this speed, significant numbers of research projects could begin to run real-time graphics applications over the network. If 5000 projects were to use the second stage network for their research, the cost per project would be $10K per year, which is a very small increment on the funds spent already for this research.

(3) The third stage is to upgrade a few selected sites to speeds of 1-3 gigabits per second (3 billion bits per second) and providing T3 service to about 1000 research sites. The operating costs are estimated to be $200M per year when the third stage is installed ten years after the beginning of the program. Again, this will be a small addition to the funds spent on the research projects which will use this network. It should be emphasized also that this network will benefit all sectors of the U.S. research and development effort. With proper planning and leadership, the costs could be shared between several government agencies and industry. Assuming 50% industrial cost sharing and the involvement of ten federal agencies, this network would require an average of $10M per year per agency.

The scientific horizons of thousands of research projects have been expanded through the principle of ubiquitous access. In practice, this means providing a supercomputer account and network connection for any researcher whose project has been judged meritorious through a peer review process. However, as the supercomputers and the connecting networks have now become saturated, many high-quality projects will not be accommodated.

It is interesting, and ironic to note that in the early days of the NSF initiative, we received a large amount of criticism for creating too much capacity. Presently, we hear that we did not create enough.

It is also interesting to note (as I did in a meeting in Austria last week) that just when the US program is starting to decline, a large number of other countries are struggling to imitate it.

The original NSF plan was to proceed in three phases over the first five years of the program:

Phase 1 was to immediately take advantage of existing supercomputer resources, such as the independent centers at Minnesota, Purdue and Colorado State, and established networks like ARPAnet and MFEnet. Phase 1 was never envisaged as a permanent solution. However, it was essential in order to get the program off the ground. This phase lasted a little over two years, and is now completed.

Phase 2 was the creation of an entirely new national computing resources, consisting of a number of national supercomputer centers connected to all research universities and laboratories through a national network, NSFnet. Originally, the number of centers was supposed to be ten to twelve, but budget limitations forced us to compromise on five. The overriding principle of the design of Phase 2 was to use existing state-of-the-art technology in order to provide a functioning system in a minimum amount of time. Phase 2 is now fully functional, and, as stated before, fully saturated. The NSF centers have in a remarkably short time, established themselves as leaders in computational research, and as showcases of new technology produced by the American supercomputer industry. NSFnet has provided connections for more than 200 universities and laboratories with the national centers and with each other. It, together with its internet connections to other research networks, has become the principal channel of communication for computational researchers everywhere in the U.S. Recently, international connections have been made, which will lead soon to a global research network.

A second design principle of the NSF initiative was that of "top-down" management. The budget was taken off the top of the NSF appropriation, rather than accumulated from the bottom. This was done because the various NSF disciplinary programs would never have been able to agree on an overall single design, but would have fractionated into separate incompatible systems.

Besides impact on research, the NSF initiative has made significant contributions to education and training. Thousands of students and other researchers have become acquainted with, and experts in, the use of advanced supercomputer and network technologies. Many of the supercomputer experts now found on US campuses received their initial training in the regular courses or special summer institutes run by the NSF centers.

Another, and unexpected, result of the NSF program, was the proliferation of "independent" centers. Recent examples are the Ohio State Center, the University of Alabama, and my own center at the University of Kentucky. These new centers, supported primarily by state and local funds, owe their existence to the strong support of faculty and students who became aware of the possibilities of supercomputing research through access to the NSF centers. This has been a very valuable "spin-off" effect which has greatly increased the capacity of US academic computational resources.

Phase 3 was envisaged as a program to advance the state-of-the-art in supercomputers and associated technologies. This would optimize the computational environment through the development of user-friendly, standardized (therefore portable) software, new numerical algorithms, advanced work stations and graphics display devices, better networks, and providing test-beds and beta-sites for new generations of supercomputers. Again, budget limitations forced us to abandon almost entirely this phase before it was properly started. A key part of the Phase 3 program was the upgrading of the centers and networks, as rapidly as the technology would allow. If budget cuts eliminate the upgrades to the next generation of supercomputers in the centers, and the next generation of networks between the centers, then the program will become rapidly obsolete, and the national supercomputer picture will revert to the pattern of pre-1984 (before the NSF initiative), when only a few well-funded national laboratories and less than a handful of universities will be able to provide state-of-the-art access for their researchers.

Some of the Phase 3 research has been, or will be picked up by other other sources, such as industries and independent centers. For example, the Center for Computational Sciences at the University of Kentucky, will concentrate much of its attention on the generation of new mathematical algorithms and the optimization of scientific applications software, which are key needs in the improvement of the supercomputing program. However, these efforts need to be coordinated at the National level, and unfortunately, NSF resources do not allow it to play that critical role.

This is essentially my second point,- that inadequate budgets will inevitably lead to obsolescence and the loss of leadership which the NSF now exerts. The momentum which the NSF has built up will quickly disappear.

The OSTP report of November 1987, "A Research and Development Strategy for High Performance Computing" maps out a plan for building on past success, and revitalizing the national program. The programs recommended by that report in hardware, software, networking and basic research, represent a significant opportunity for federal stimulation of supercomputer-related research. Supercomputing is too important to the nation's competitiveness position to take a laissez-faire attitude. Certainly, the governments of competing countries have already embarked on initiatives at least as large as proposed by the OSTP report.

My third point relates to other funding mechanisms which have been suggested to solve the budget dilemma. However, they are in my opinion unsatisfactory, in that they would tend to distort the nature of the program in undesirable ways.

One mechanism which has been much discussed is to charge the end user for the use of the supercomputer centers and the networks. Inevitably, this would lead to a demise of the national program. There is no way that thousands of individual grant budgets can be conglomerated for a single unified program. No other national research facility (such as a telescope or particle accelerator) has a charge-back mechanism, since it is well-known that they cannot function at the state-of-the-art when a thousand funding sources have control over budgetary planning. Local cost-sharing is, on the other hand, desirable, in order to fund local area networks, gateways and specialized computers for processing supercomputer output, but the bulk of the budget (and therefore the planning) should come from a central source.

Other funding agencies might fill the gap, but there must be a central planning process. It will be very difficult to get the various mission-oriented agencies to line up their priorities along a unified direction. There should be a lead agency, which is most logically the NSF, since it has the broadest charter. NSF puts the bulk of its resources into academic research, which is the source of a large part of the creative ideas in fundamental research in this country, whereas the mission agencies focus their funds on their own laboratories. It is important, of course, to include the national laboratories in a national program, but they should not be the only focus.

Industrial laboratories, as well, should be a part of the program. Some of the NSF program already benefits from their contributions. But again, their priorities would distort the direction if they were to take over the major part of the planning.

To summarize this point, ideally there should be a lead agency for the National Computer Initiative, and this agency should be the NSF, since it has already demonstrated its ability to develop a broad based national computing facility, and is the natural conduit for the creative energies of the academic community.

Alternatively, if the notion of a lead agency is not found acceptable, my second choice would be to establish a central planning process managed at the highest level of the Executive branch (for example, OSTP). This should involve full-time staff to coordinate with designated liaisons in the relevant agencies, in the academic community, and in national and industrial laboratories. The central planning staff for the National Computing Initiative would work closely with OMB and Congress to work out budget priorities, and to determine the special add-ons for each agency's budget for their respective components of the overall plan. The charter of this structure would be broader, and have more responsibility than the current FCCSET committees.

Of prime importance in the National Computing Initiative should be the National Network, of sufficient capacity to allow convenient remote access for any meritorious project. The basic principle should be to provide open access with quality controls. Most research facilities (including the NSF Supercomputer centers) have a review process to ensure fair and equitable division of their resources. Similar controls need to be worked out for the the national network. (Most networks, including NSFnet, currently have no control over how much access an individual user can have.)

The internet concept is equally important, so that the (eventually global) network becomes as easily accessible as the telephone, available to everyone who merits it.

State-of-the-art computational resources (supercomputers and beyond) will naturally connect into the network, no matter where they are, or who funds them.

The only problem of building the network is achieving increased capacity, and the technology is very close to the gigabaud rates required for future applications. All that seems to be required is the will to do it. The success of the venture is virtually guaranteed, based on our extrapolations from past experience. It is also generally accepted that such a supernetwork will be the major communication medium for the knowledge-based (or information) society which is becoming dominant today.

The implications of the National Computing Initiative are significant for all future progress in research, development and technology, and I strongly recommend that this committee support it.

QUESTIONS OF THE CHAIRMAN AND THE ANSWERS THERETO

Q1. There is an inherent tension between computer scientists who want to have a high-speed network to experiment with and the computational scientists (Physicists, engineers, and others) who need networks to access supercomputers. Computer scientists who study networks want state-of-the-art technology. Other scientists may not need a state-of-the-art network, they just want it to be reliable.

- Is it possible to resolve this conflict?
- Do you think that NSF is doing a good job of keeping both communities happy?

A: You are quite right in your analysis of the competition between computer science and scientific computing (also known as computational science). Computer science seeks to advance the state-of-the-art of computers and networks, where as the scientific computing community needs computers and networks to solve physics and engineering problems. These different priorities leads to slightly different funding strategies. Computer scientists want to concentrate on novel, experimental approaches, whereas the physicists, engineers etc. want more established (but still state-of-the-art) and robust instrumentation.

Since the terms "computer science" and "scientific computing" are often confused in the funding agencies, there is an unfortunate competition for funds which sometimes becomes quite heated. What is needed is for the funding agencies, like NSF, to distinguish clearly between the two strategies and to maintain a reasonable balance in the budgets for the two communities. Also, there should be an effort to reconcile the differences between the two communities, by establishing the principle that their long-term goals are the same. If this is accepted, a common strategy might follow, in which there would be a natural transition from the experimental computers and networks to the user-friendly and reliable computers and networks required by the scientific computing community.

The third stage network will require the investment of some development funds, estimated to be $40M per year for ten years. Under the same assumptions of centralized planning, industrial cost-sharing and multi-agency involvement at the federal level, this works out to an average of $2M per year per agency. (Of course, those agencies which are now providing the leadership in networking, such as NSF, DARPA, DOE and NASA, would presumably need to be funded at levels higher than this amount.)

To summarize the previous points, the OSTP cost-estimate of developing the gigabit network is reasonable, and a modest increment on existing funding. Certainly, these costs will be a small percentage of the amount which will be invested in total computing capacity over the next ten years. The major uncertainty is in the time schedule, but I believe the OSTP estimate is on the conservative side. With appropriate leadership, the gigabit network could be achieved in less than ten years. I understand, for example, that a pilot project at this speed is planned to be operational within three years.

At this point, I would like to reemphasize a major point. The OSTP cost-estimates are based on the assumption of centralized planning and leadership. If each federal agency were to develop its plans independently, the costs could easily be ten times as high.

Q3: We've heard a lot about how supercomputers can produce billions of bits of data a second. And the proposed National Research Network will transmit billions of bits per second.

- But how do we store these incredible amounts of data?

- Are we paying enough attention to the question of data storage?

A: You are quite correct in pinpointing mass storage requirements as a critical element in the national computer initiative. At the present time, because of cost, magnetic disk space is extremely limited at a typical supercomputer center, perhaps a few hundred gigabytes at the most. This is a severe limitation for many applications, especially those which rely heavily on graphics. Supercomputer-generated images can easily saturate a center's capacity. Usually, a user must transfer such data to magnetic tape or cartridges which are economic but are slow to generate and access.

The hope of the future is optical storage. At the present time, however, optical technology which can provide rapid, cost-effective read-write capability is not available. Certainly, there is a great deal of industrial activity in this area. But, as far as I know, there is no federal program to encourage the use of this technology for scientific applications. If we are going to take full advantage of the next generation of supercomputers and networks, such programs are going to be essential. This, unfortunately, was not emphasized in the OSTP report, and any federal plan should include a significant amount of support for providing state-of-the-art storage technology to the centers and research projects which require it.

Senator GORE. Very good, thank you very much.

Our next witness is Dr. William Wulf at NSF.

Welcome, please proceed.

STATEMENT OF DR. WILLIAM WULF, ASSISTANT DIRECTOR, DIRECTORATE FOR COMPUTER AND INFORMATION SCIENCE AND ENGINEERING, NATIONAL SCIENCE FOUNDATION

Dr. WULF. Thank you, Senator. As you said, I am from NSF. I am relatively new; I've only been there a couple of months. I am the Assistant Director for Computer and Information Science and Engineering Directorate, which, among other things, has responsibility for the supercomputer and, the networking programs, as well as for basic research in computer science.

I want to thank you for your interest and actions in the past, especially in the area of computer networking. I have a deep personal commitment to the importance of the computational infrastructure for science and ultimately for the country. It is NSF's actions in this area which really attracted me to come to the Foundation. It is just so terribly important.

I recognize the time so I will highlight just three major points if I may.

First of all, I want to stress the importance of the computational infrastructure to productivity, competitiveness, and national security. I think we all recognize that the conduct of science is changing, and that it is changing rapidly. I am not sure we all appreciate what that means.

The infrastructure has an enormous impact on scientific productivity, I think we all recognize that even if we cannot measure it very well.

It also has profound impact, on the pool of scientists available to work on any given problem. In particular, the network opens up the possibility for researchers at four-year schools and minority institutions to work on important research problems, to participate as would anyone at a major research university.

The network also enables previously undoable research—research that simply would not have happened because two persons could not have collaborated because they did not have access to necessary instrumentation.

Finally, it has enormous impact on technology transfer. One of the phenomena that I observe in the computer industry is that academics and industry researchers and product developers have become part of the same intellectual "stew." Ideas are transferred more quickly in both directions. Important problems to industry are also exposed to the academic community, and academic ideas and utilized much more rapidly because all of the relevant people are connected to network.

The second major point is one that several people have mentioned, but I would like to reiterate. It is the need for balance in this program, balance between the four things mentioned in the FCCSET report; high performance computing, networking, software and algorithms, and the human infrastructure.

It is far too easy in a political atmosphere and in our excitement over the hardware to overlook the importance of software and people. We know in our supercomputer centers, for example, that there is a potential for something like at least a twenty-fold increase in performance by using better algorithms, and better trained people. It is not going to be easy to realize that increase; basic research is required and it is extremely important.

You mentioned earlier that there is a lack of focus evident in the software area and I think you are absolutely correct. The software problem has been identified for the last 20 years; we have tried numerous quick fixes, but they have not worked. The problem is still with us. It is time to recognize that it is a very difficult problem and that basic research is needed. It is going to take a while to solve it.

Senator **Gore**. Well, I think an institutional fix is needed, too, to single out the applications that are most critical.

Dr. **Wulf**. I agree. One of the tasks I would like to see added to those of the supercomputer centers is to become intellectual centers, among other things, in the development of software and algorithms.

The third point is that basic research is required in all of the areas that we have talked about. Phase I and Phase II of the network require development, require management, but Phase III requires a great deal of research, not only in the physical devices, the switching machines and so on, but also in software provided services to make it a useful entity.

I would like to point out that in my directorate we are being as aggressive as our budgets allow in following up on the provisions of the FCCSET report. We have created a division, Networking and Communications Research and Infrastructure and, based on the Appropriations Bill which has just passed, I think we should be able to increase its budget by about 50 percent between 1988 and 1989. We have the 1.5-megabit NSFNET back bone up and working. In fact, the use of fiber envisioned by your original bill is here, it is working now.

We did it, by the way, on time and at budget, and we did it with State, Federal and industry collaboration, that is a model of what needs to be done throughout this area.

We will pursue developing the computational infrastructure as vigorously as we can, given the budget constraints. As far as I am concerned, that infrastructure is my highest priority at the Foundation. It does need to be a balanced infrastructure. It is not just a matter of buying hardware.

There are many challenges in getting there—intellectual challenges, organizational challenges, managerial and political ones—but the payoff is enormous. You said before we do not know how to measure the payoff; that is correct. On the other hand, we spend $60 to $70 billion a year of Federal money on basic research. Even a 1 percent increase in productivity on that research would pay for much of the cost projected by the FCCSET report.

I personally believe that the productivity increase will be measured in intregal factors, not in small percentages.

Senator **Gore**. Clearly.

Dr. **Wulf**. When you add to that the leverage that research provides in terms of the gross national product, it is just mind-boggling.

Senator GORE. Ours and that of other nations?

Dr. WULF. Indeed. Anyway, I look forward to the opportunity of working with you for the next couple of years to create the computational infrastructure.

[The statement follows:]

STATEMENT OF DR. WILLIAM WULF, ASSISTANT DIRECTOR, COMPUTER AND INFORMATION SCIENCE AND ENGINEERING, NSF

Introduction

My name is William Wulf. I am Assistant Director of the Computer and Information Science and Engineering (CISE) Directorate at the National Science Foundation. I came to NSF from the University of Virginia, where I hold the AT&T Chair in Engineering and Applied Science. Prior to that, I was Chief Executive Officer of Tartan Laboratories, a software company I helped found. Before that, I was Professor of Computer Science at Carnegie-Mellon University for 13 years.

I am here today in several capacities: as tne NSF Assistant Director responsible for NSF networking activities, the supercomputer centers, and computing research; as the current Chair of the Networking Subcommittee of the FCCSET Committee on Computer Research and Applications; and, less officially, as an interested scientist with experience with research networking applications, and a user of the networks since the early 1970's.

I will focus my remarks on networking, but also comment on the critical interaction of networking and computing, and the need for a stronger underlying investment in software and human resources. The FCCSET report, "A Research and Development Strategy for High Performance Computing," (November, 1987) outlined a balanced program leading to a national research infrastructure consisting of four components: high performance computing resources (supercomputers and local computing), a national research network, software technology, and a strong human resource base.

In the excitement over the hardware, software and people are too often overlooked. Our own experience with supercomputer support suggests that is a costly mistake. For example, with better algorithms and "tuning" the software, there is the potential to increase the performance of the machines currently at the NSF Centers by up to a factor of 20. With the next generation of parallel supercomputers, the investment in basic software research and training people in the new technology will be even more critical to the overall efficiency of the program, and hence to the quality of the science which emerges.

NSFNET

NSFNET has become the core of a rapidly developing "internet" which ties together many governmental, private and international research networks. Because of our unique mission, NSF has become the de facto lead agency in research networking, and advanced networking research has become a very rapidly growing part of CISE. The new 1.5 megabit/second NSFNET backbone became operational last month, and is working well. This is one step, but a critical one, in the plan outlined in the FCCSET report. It is also testimony to what can be accomplished with Federal, State, and industry cooperation, and the skills of university based organizations like Merit, Inc. The State of Michigan, IBM and MCI deserve much of the credit for the rapid implementation of the initial part the "Phase 1" network recommended in the FCCSET report.

Research Networking: Prospects and Benefits
The agencies represented on the FCCSET Committee recognize the importance of information processing and communication to research productivity, and consequently to our national competitiveness, security and prestige. At NSF, for example, that recognition has been translated into concrete terms. We have steadily increased our support for networking over the past several years, and expect to continue the trend within our overall approved budget totals.

These are exciting times in computer networking, in high performance computing, and in information processing in general. We are at a stage in networking analogous to the 1890's in voice telephony, or the 1930's in interstate highways. We have small, narrow markets but interesting economic and societal possibilities if we can overcome the multitude of separate islands of connectivity, develop the "right" services, expand the directory that makes it possible to find one's way through the system, and establish standards.

There are problems to be solved to realize the potential of networking, but experience clearly suggests that the payoff to the nation will be substantial. There are a great many issues to discuss, but first I want to concentrate on just a few of the benefits of networking as I see them.

It is clear that networks increase the productivity of researchers, by allowing them to collaborate with a larger set of colleagues, by giving them timely access to data collected at remote sites, and by giving them access to unique resources. Networks can also effectively increase the pool of researchers, for example, by making it possible for the faculty of four-year and minority institutions to become full participants in research projects based elsewhere.

More important than either of the above, however, is that networks allow research to be undertaken that simply couldn't be done otherwise. For example, they permit interdisciplinary collaboration between researchers that do not happen to be neighbors. They permit access to data and facilities that would not be accessible otherwise. They speed the spread of ideas, broadening the scope of problems undertaken by the research community. They aid technology transfer by allowing direct and fast interaction between people, and between people and research literature.

We cannot quantitatively measure these effects, and have no way to predict their future implications. However, I believe that what will emerge from exploiting more capable, less expensive computing and communications technology, and vastly more information in computable form, will have profound effects on science, engineering, and society. One can see glimmerings of this future in the major successes research networking has already had, for example, in medical informatics (SUMEX-

AIM and National Library of Medicine's Medline), exploitation of VLSI technology design and prototyping services (MOSIS), and the NSF supercomputing effort. I have personal experience with the design of Ada, the new DoD computer language for embedded systems, where the involved community was located literally throughout the world, and where, I believe, the work simply could not have been done without ARPAnet.

The Federal investment in research exceeds $60 billion each year. Even a small productivity improvement will have huge effects - $600 million worth for each percentage point gain. I believe that the effect will be measured in integral factors, not in small percentages!

Networking increases the effectiveness of the federal research agencies. Again, we cannot quantitatively measure the impact, but many of the same effects that one observes among researchers are also evident between the Federal research agencies and the research community.

For example, DARPA has been able to coordinate new initiatives, and to rapidly communicate findings among research laboratories and industry. NSF is reviewing proposals more quickly than with conventional mail, has improved the process of interacting with the research performing community, and is experimenting with electronic proposal submission and automated connectivity to the administrative and decisions support systems. DOE contractors are sharing libraries of supercomputer code to use the computing power and expertise at the national laboratories more efficiently and effectively. NASA is exploring the potential of "telescience" to more effectively utilize remote facilities such as the space telescope. NIH is making the new genetics data bases more available, more quickly, to medical and biological researchers.

National R&D Partnership and Competitiveness

Increasing the effectiveness of researchers is important, but the real payoff is when the results of research are applied. Networking has a critical role to play here too, in moving ideas more rapidly to the private sector, and in generally increasing the "intellectual stew" from which our industries extract the ideas that are ripe for commercialization. Moving people is still an important way to move ideas, but increasingly, industrial researchers and product developers are becoming a part of the same electronic community as the academic researchers. The two-way flow of ideas in this enlarged community is not only benefitting industrial developers, but is providing stimulating, realistic problems and critique to the academics.

NSF's extensive experience with a predecessor to NSFNET, called Computer Science Net (CSNET) illustrates the importance of connectivity to industry research laboratories. Industry has invested heavily in internal corporate networking. Those joining CSNET paid at a rate which helped subsidize the university participants. They believed it was worth the price.

Similarly, industry has invested heavily as partners in NSF networking activities. Apart from IBM and MCI financial participation in NSFNET, there has been substantial investment by Digital, Apollo, Apple, IBM, and SUN in the EXPRES project, which explores advanced documentation and collaboration technology. This industry-government partnership results not only from industrial understanding about the long term payoffs from basic research, but from the historic, enlightened partnership of American industry, government and universities.

American educational institutions have a history of leadership in developing and using computing and information technology. They pioneered network-based computing innovations such as timesharing systems, research software packages, standard operating systems, automated text retrieval, and computer aided instruction. BITNET was created to enhance scholarly communications, and innumerable instructional applications have been developed and placed into routine use.

NSF has been working closely with EDUCOM and other educational consortia to make NSFNET a reality, and to insure that it serves the needs of teachers as well as researchers. We will continue to do so, because the need for improved access to educational resources parallels that for research resources, and because the implications for improving educational productivity and saving scarce resources are profound.

Networking Challenges

The potential benefits of networking will not be realized without meeting a number of challenges. The U.S. is not alone, and perhaps is not even the most aggressive, in recognizing the benefits and addressing the challenges of information technology. There are several types of challenges.

First, there are intellectual challenges. Networking advances require innovations apart from the physical net technology. These include the software used to control the net, the services (including directories, addressing and routing, access and security controls, libraries, automated retrieval and analysis utilities, tools for collaboration, resource sharing and interaction), and management frameworks and tools for tracking the use, utility, and effects of the network. Research in these areas is as important and as intellectually challenging as any for advancing understanding and promoting technology transfer. The challenge is to bring together the many disciplines and their specialized tools to work together on the problems involved.

Second, there are substantial organizational, and managerial, challenges at all levels, which make it difficult to generate the consensus needed for concerted and sustained action. The problem is real and is exacerbated by the fact, as pointed out in the NAS critique of the FCCSET report, "Toward a National Research Network", that much of the federal, state, industrial, and university planning must be done without the benefit of a firm specification. The challenge is to build the coordinating mechanisms for the creation and management of the network.

Third, there are exciting technical issues to be resolved. The "phase 1" and "phase 2" networks envisioned in the FCCSET report are reachable using today's advanced technology. Advances in switching technology, communications protocols, routing, and accounting algorithms are necessary to realize further progress. A well coordinated effort by government, industry, and universities will be critical to achieving theses advances.

Fourth, of course, there are financial challenges. Mechanisms must be created to sustain the unusual mixture of research and service roles for the network, to deal with the costing and accounting, recovery and subsidy problems of a system with central and user controlled segments. These might seem overwhelming at first glance, but we are dealing with a national resource which will be far less costly if developed in a coordinated way than would be the case if the parties at interest continue on separate paths .

Finally, there is the challenge of maintaining the balanced program of computing, networking, software, and human resource development, advanced in the FCCSET report. Although there are many analogies between the information processing infrastructure we are developing and the familiar infrastructures of the past such as the highway system, there are also important differences. Specifically, the information processing infrastructure is not just physical devices composed of glass fibers and silicon chips. The physical aspects of the infrastructure are easier for the lay person to relate to, to point to, and to be proud of -- but without software to activate the hardware, and trained people to use it, the gadgets are literally meaningless.

Conclusion

The near term payoff of the new information technology infrastructure is large, but the longer term impact is *far, far* larger. We are privileged to be involved in the development of this infrastructure. A few centuries hence, this time will appear a milestone to human progress. We are developing tools that embody our knowledge and amplify our intellectual prowess. The industrial revolution, by contrast, was the culmination of experience with tools that amplify our physical prowess.

The easy-to-foresee payoffs of the effort today will be innovation important to R&D productivity and industrial competitiveness. Yet, since our intellectual capacity is the essence of what makes us human; tools to amplify that capacity cannot but have an impact beyond our ability to imagine.

Our severest challenges, then, are to support the visionaries, to catalyze the human resources to realize the potential of high performance computing and communications technologies; to accept the risks and commit financial resources to seed these developments; and to maintain the commitment and cooperation of federal agencies, private industries, state governments, universities, and most of all, the individuals who have to be involved to make it work.

Senator GORE. Thank you very much. I look forward to it also.

Our last witness today is Dr. Alan McAdams, Associate Professor at the Johnson Graduate School of Management at Cornell.

Welcome. Please proceed.

STATEMENT OF DR. ALAN McADAMS, ASSOCIATE PROFESSOR, JOHNSON GRADUATE SCHOOL OF MANAGEMENT, CORNELL UNIVERSITY

Dr. MCADAMS. It is a privilege to be here today. I would like to point out despite the fact I am an economist, I am a member of the IEEE and I serve on its Scientific Supercomputer Subcommittee, on its Technology Leadership Subcommittee and on its Committee for International Competitiveness.

My comments are consistent with many others you have heard today but I want to express in slightly different words that what we really need is a systems approach. Systems approach is required.

Dr. Wulf mentioned balance. I think that is one element of the systems approach. We need a systems approach to the national education research network. We also need overall technology policy. Those two are interrelated. I think I will show these relationships in just a moment.

One further illustration of the relationships comes because of the success to date of our backbone, and its new backbone really comes from the fact that we have something on the other end of those backbones. We have state of the art supercomputer systems. The backbone alone is not very useful.

We need to interconnect supercomputer centers that stay at the state of the art. That is going to be a big order. They are now, but the art is moving quickly.

We need a production network and we also need a test bed network, that dual approach to the network will permit the constant upgrading to later technologies.

We need to have flexibility built into the way we prepare a network. We need funding that assures that we have an integrated system from end to end. It does not do much good to get from one end of the backbone to the other end of the backbone if you happen to be a researcher in a university not on the backbone.

The way funding takes place will influence the opportunity to have an integrated system. We have got to be able to find something in the system, we need to be able to correct errors quickly and easily.

Let me put in a comment of regional pride. One of the regional networks I know about is NYSERNET. For a year I have been working on research for those people and NYSERNET today is at T-1. It is one and a half million bits now and my understanding is that NSFNET will move users to that speed on January 1. At the moment a user has access to only one-third of that rate.

NYSERNET in the National Net '88 presentation that Mr. King referred to and was proposed to be a two-stage network. This was proposed by people from New York Telephone Company and NYNEX with a production network and a test bed network.

That seems to be a very good model that would be quite effective.

I would like to refer to the paper that my colleagues and I presented at National Net '88 and it has some relevance to the comments made by others.

First, I think we have to recognize there are two kinds of information. One of them is new information, the type that is appropriate for use on a research and education network, and the other is current information, it manages the economy as it stands. We regard data network, a data network devoted to research and education, as absolutely critical to the economic future of the United States.

And we show how it differs from the commercial-oriented networks. The bulk of its benefits accrue to society as a whole and not to particular organizations or individuals that create them.

Since these public benefits cannot be sold in a market, research and education networks cannot be self-generating or self-financing.

Government funding thus is both essential for the survival of these networks and fully consistent with the rationale for the free market. It is a means of fulfilling a gap in the ability of the market to bring such benefits into being.

A number of people have said they want to have researchers have a choice. By giving them money they can choose among networks. Where does the money come from? The money for researchers must come from government, so whichever way you are doing it, you are really saying that the money is coming from government for the R&D portion of such a network.

We also recognize that this is consistent with other national initiatives such as the mail system, continental railroads and the one which we have been referring to often here, the highway network. I would like to point out, however, there is no commercial highway network that was in competition with the national highway network like there is in the instance we are talking about, that raises some problems.

The network will come about. The real question we have is, How and when? We cannot afford to wait until the United States is pushed into an effort that is either too little or too late. The time to do it is now.

Now, back to the systems approach. If the only thing we do is to provide research and education network, that could be counterproductive. You have mentioned it before today. It further opens the U.S. market and technology base and unless we can outperform others in bringing research to the level of innovation and productivity, which is then required for maintaining differential standard of living, we will lose out.

So if we go just with this network and do not take care of the full system such that we can respond to what the network generates, we will lose more quickly than we will win.

Now, the IEEE has just issued a paper on the vulnerability of our supercomputer base. I worked on that paper and I think it is important to point out here because it shows that that vulnerability is a symptom of our current problems and overall technology, having a wonderful education and research network and not taking care of problems where technology can further open our system to exploitation by others.

What we perceive in the IEEE statement is that the civilian technology industry and trade policy requires coordination to be effective. It can only be provided at the national level. So what people have said is only at this level will it be possible for a true systems approach to be taken to the problems facing the country in international competitiveness.

Action is required on long-term R&D policies. It is necessary to integrate the activities of government industry and universities, it is necessary to achieve timely, effective technology transfer from stage to stage in the process. This is most important with the network. It is necessary to ensure careful attention to the provision of suitable people power, something others have just mentioned, at each stage of the process. And that involves not just numbers but appropriate intellectual capital.

We have talked about institutional framework for doing this. The challenge is to find an acceptable institutional framework in which government, industry and academia can pursue the objectives for the long-run benefit of the nation as a whole. A necessary requirement of that framework is that it must ensure economic and technological decisions are taken on the basis of economic and technological as opposed to mainly political or military criteria.

It may well require coordination and leadership functions be vested in a new, lean, expert civilian agency of government that is capable of focusing on longer-term national interests.

Only through a coordinated approach to all the above issues will we be able to ensure a strong U.S. base for innovation productivity and international competitiveness, a base that includes supercomputers as a vital factor and most importantly to which a national education and research network must make a separate contribution.

[The statement and material referred to follows:]

TESTIMONY FOR PROFESSOR ALAN K. MCADAMS
Johnson Graduate School of Management, Cornell University

"Economic and Management Factors in a National Research and Education Network,"

Thank you for inviting me to appear today. It is indeed an honor to appear before you, especially in such distinguished company.

I have been asked to testify on several aspects about a national research and education network. My comments deal with the economic rationale for federal funding of such a network, the role of a national research and education network in U.S. international competitiveness and the current status of research and education networks in the U.S. In an accompanying paper, my colleagues and I attempt to make clear the economic rationale for federal funding for a national research and education network and the important role that such a network can play in U.S. international competitiveness. In that paper we also explain why these conclusions are consistent with a market-based economy. We also suggest that the costs of creating such a network are likely to be modest in relation to other important national programs. We conclude that a national research and education network is a necessary element for U.S. international competitiveness.

In that paper, "Economic Benefits and Public Support of a National Research and Education Network", (preprint attached)[1] we identify a dichotomy between "new" information and "current" information. "New" information is that information that brings about change, including fundamental change, in a society. "Current" information is the information used to manage the society in its on-going form. It is the "new" information on which innovation is based, and innovation is the key to our international competitiveness.

Information of this kind is likely to have enormous "spill-over" benefits to parties not directly associated with its creation. That is, the benefits from the change induced by "new" information spill-over to the society at large. As a rule, benefits of this kind are difficult if not impossible for the parties creating the information to capture. Thus it is virtually impossible to create a market for the allocation of such information. In the presence of such a "market failure" and in the absence of private incentives for creating such information, government action to ensure that its benefits are brought into being is both appropriate and consistent with a market-based economy. In our paper we establish that a national research and education network--a network devoted to facilitating the development of "new" information--is fully appropriate for federal funding and support and that it can be an important contributor to U.S. international competitiveness.

---

[1] The attachment was not reproducible.

Today I want to emphasize that the economic rationale for federal support of a research and education network extends to all those elements in the network which produce significant spill-over effects--effects that cannot be captured privately in a market. This implies that federal support is appropriate not merely to a backbone network, but also to those regional and local networks which provide spill-over benefits to society. It is the presence or absence of significant spill-over benefits that establishes the basis for government funding, not an arbitrary characterization of a network as "backbone", "regional", or "local".

Given that resources are limited, however, we recognize that rationing must take place. It is for this reason that we suggest two qualifications. First, that federal funding should be limited to networks supporting "new" information with the greatest spill-over benefits; and, second, that it is appropriate for state governments to share the funding for education-related activities, an area long recognized as a province of the states.

A National Research and Education Network: An Element in an Integrated National Technology Strategy

Having made statements about the desirability of a national research and education network and the need for government support to bring it about, it is necessary to put those statements in context. We, along with many others, believe that the U.S. high technology base is fundamentally threatened as long as the U.S. continues to follow its current policies. While a national research and education network is an important, necessary component for a resurgence in U.S. technology, it is far from sufficient. The accompanying paper, "U.S. Supercomputer Vulnerability", prepared by the Scientific Supercomputer Subcommittee of the Committee on Communications and Information Policy of the IEEE (in the preparation of which I participated) and issued by its U. S. Activities Board, illustrates serious problems in the U.S. technology base, their depth and pervasiveness. (Perhaps this should be the topic of another hearing.)

Unless the national research and education network is incorporated into a system-wide solution to the properly-diagnosed problems facing U.S. high technology, it could do as much harm as good. Swift, accurate information flow is a two edge sword; it can permit the rapid imitation of even the most advanced U.S. innovations, as well as speed the pace of creation of such innovations. In other words, the downside of an excellent networking system is that it can facilitate the rapid transfer of technological innovations to other countries with lower standards of living where they may be exploited at lower costs. We need to speed our rate of technological innovation. At the same time we need to maintain our technological leads for periods sufficient to sustain our differential standard of living.

Objectives to be Sought

The major objectives of the national research and education network should be to provide effective:

- connectivity,
- transparency,
- speed.

All three are required to permit rapid intercommunication among researchers and educators. The meaning of the term "speed" is obvious, but the other terms require some explanation. By effective "connectivity" we imply the ability for one research person to reach any other in the country, and overseas if necessary. By "transparency" we imply the ability of that researcher to reach any other through a network whose functioning is so smooth and automatic that the user does not even know it is there. Another way to characterize "transparency" is to say that the user, when accessing resources through the network, experiences no more difficulty than when that user is operating his or her desktop computer system of whatever make (e.g., Macintosh, PC) as a stand-alone work station.

Current Networks

The world we live in today is a far cry from one in which a high speed research and education network provides universal connectivity and true transparency to the user. It is closer to a tower of babble than it is to such a network (see Appendix). Let me illustrate with a few comments about three networks. The three are:

- ARPANET
- BITNET
- NSFNET

You have already heard many comments about these networks. ARPANET, created by DARPA (the Defense Advanced Research Projects Agency), is the "grand-daddy" network that proved the feasibility and efficacy of the networking concept. (As you know, it is being systematically dismantled today due to a lack of funding and the advance of technology.) In any case, it has been a victim of its own success. It became heavily congested to the point that only transmissions of limited size could be expected to be delivered through it. Its bandwidth never exceeded 56K. BITNET is a cooperative network currently maintained by a group of universities each of which has assumed responsibility for a host system plus a tied-telephone line to an adjacent host. BITNET is approaching an ubiquitous status in the U.S. and connects to many subnets overseas. The bandwidth of BITNET is 9.6K. It is used mainly for electronic mail. It represents only the most limited of file transfer capability. It is approaching the limits of its usefulness, in part because of congestion. Initially the highest level of connectivity was achieved through ARPANET, the focus of which was on defense related missions. High connectivity is achieved today among researchers at universities through BITNET, but because it uses proprietary protocols, it is difficult to interconnect to other networks..

Neither of these networks can be said to have aspired to transparency for the user. Rather the opposite has been the case. The users have found it necessary to be excruciatingly aware of arcane address structures and command sequences in order to make use of these networks. Often, as has just occurred in the case of the main host link for the U.K., address changes are made by the host institution without notification to users. In this instance, no messages were able to get through from the U.S. to the U.K. beyond the Computer Science Department of the University of London because the main U.K. host link was changed-over to serving only that Department, while different hosts were assigned to transfer traffic to other sites in the U.K. and Europe--with no advanced notification to users. (A flurry of telephone traffic across the Atlantic ensued to attempt to diagnose the problem and then to re-address transmissions.) No mechanism exists for avoiding such problems or rectifying them quickly or easily once they have occurred.

NSFNET is in a period of transition. Until July 1, 1988, it was also limited to a 56K bandwidth. On that date it changed over to a T1 network initially operating at approximately one-third the T1 rate, or 440K. On January 1, 1989, it is scheduled to go to full T1, or 1,544K. As has been stated, its basic mission is to provide a high-speed backbone between the NSF-sponsored supercomputer centers throughout the country. The new backbone represents a significant step forward in this regard. Not only does it connect the five NSF supercomputer centers, but it also provides direct connections to seven regional networks. As you also know, the five NSF supercomputer centers are located at:

- U.C. San Diego (Cray)
- U. Illinois (Cray)
- Pittsburgh Consortium (Cray)
- Princeton (ETA)
- Cornell (IBM)

In addition, NSFNET connects to NCAR at the University of Colorado plus seven regional nodes:

- SURANET @ U. of Maryland
- MIDNET @ Nebraska U.
- MERIT @ U. of Michigan
- WESTNET @ U. of Utah
- SESQUINET @ Rice U.
- BARRNET @ Stanford U.
- NORTHWESTNET @ Boeing Computer

(NYSERNet, a regional net about which more is said below, connects at Cornell, one of the supercomputer sites. The same is true of supercomputer related subnets at other sites.) The old backbone nodes have been disconnected. The thirteen nodes of the new backbone have been activated and are performing smoothly. Prior to the changeover, all regional networks connected into the supercomputer centers and achieved access to the backbone through those centers.

The connectivity provided by the new backbone, while greater than for the old, is still quite limited. No attempt has been made to provide transparency through the network to the supercomputers of the various manufacturers. A user still must become acquainted with the idiosyncrasies required for access to each particular supercomputer center.

The regional network with which I am most familiar is NYSERNet, the New York State Education and Research Network. To my knowledge, NYSERNet is the only network (including the backbone) which today is moving a substantial portion of its traffic at T1 (1,544K) speeds. It is an experimental network that was developed in close cooperation with New York Telephone Company (NYTEL) and its parent, NYNEX. NYSERNet connects fifteen major research universities and the Brookhaven National Laboratory plus several industrial research sites to the Cornell National Computer Facility (CNCF). It is in the second year of a two year trial. It also has been quite successful. Like the NSFNET, it is a TCP/IP network. At the recent conference, "NATIONAL NET '88: Perspectives on the Need for Public and Private Initiatives", NYNEX and NYTEL representatives suggested the desirability of an evolution in NYSERNet into two parts: a production network providing immediate service to users, and a developmental network to push back the frontiers of networking technology. It is anticipated that both networks would evolve rapidly with new technologies cascaded from the development network to the production network once their technical and economic feasibility had been established. This model suggested for NYSERNet, appears to be highly appropriate also to a national research and education network.

A true systems approach to a national research and education network must provide for timely transitions to new levels of technology. To this end it is also essential that, as in the proposal for NYSERNet, a companion high-performance test-bed be provided in addition to a production network for research and education. Networking technology is in its infancy. Rapid, pervasive changes in technology must be planned for. Each new iteration of technology is likely to require the development of new processes and procedures. It will be necessary to provide for development of, evolution to and testing for future networks.

The Supercomputer Centers Themselves

Given the perview of this Subcommittee, it is important here to provide a further observation. A number of experts perceive that the NSF supercomputer centers have been successful far beyond initial expectations. By making the most advanced computing power available to the leading, expert researchers in our nation's universities, a far greater increment in research output has been achieved than had been anticipated. This has been true despite a number of problems which had to be overcome. The NSF experiment in providing the highest computational capability to the country's leading university scholars, has been a great success. This illustrates a further point consistent with several opening comments. A high performance national research and education network is an important, necessary element for U.S. international competitiveness, but a network is no more useful than the facilities it interconnects. These, too, must be maintained at a leading-edge capability.

Current NSF Networking Strategy

As you know, our current research and education network is a three tiered network. It involves:

- the NSF backbone
- the regional networks
- local networks at the user's site

Problems are inherent in this structure and they are compounded by the funding decisions that are being made under current budget stringencies. Though connectivity is enhanced in the short run by such a structure, transparency is virtually foreclosed. The research and education network of the future is going to have to be focused heavily on graphics and real-time video transmissions. These are virtually impossible to transmit through a decentralized, three-tiered structure. To provide perspective for recognizing the significance of the characteristics of today's networks, it is instructive to look to the future.

Many expert observers recognize that in the immediate future, perhaps within as little as five years, workstations with the power of today's Cray I Supercomputer will be on the desktops of literally thousands of users. Experience has shown a phenomenon which some consider paradoxical. When increased computing power is immediately available to users at their local workstations, the result is an increase in demand for networking capability, not a decrease. In the last week of July of this year, I had the opportunity to visit the Los Alamos National Laboratory as a member of the IEEE's Scientific Supercomputer Subcommittee. That visit cemented our understanding that a nascent, vast demand for networking capability exists that will be tappable in the immediate future. Today's supercomputers generate information at such a rate that it can only be effectively analyzed through graphic and other visual means. Los Alamos is already providing some users with the capability for the equivalent of real-time video representations of the output of their supercomputer computations. Within two years, Los Alamos anticipates making this capability available to all 8000 users of this, the world's largest installation of supercomputers. (Today four of Los Alamos' complement of supercomputers are Cray I's. They are about to be replaced by a smaller number of Cray X-MP systems which within a year will double the computational power of this installation to the equivalent of 30 Cray I's.)

Similar transitions to graphic and real-time video presentations of complex data sets are taking place at supercomputer centers throughout the country. Such graphic displays are extremely bandwidth-intensive. Researchers are seeking "hypermedia" presentations of information and data. That is, they seek to interact with each other and with their research projects through simultaneous video, voice and data presentations. These, too, are extremely bandwidth-intensive.

Los Alamos leads a consortium of manufacturers and users that have banded together to make possible standardized approaches to the problem of providing such network services. With the equivalent of Cray I's on thousands of desktops, enormous demands will be placed on the national research and education network. Once acquainted with the power of hypermedia interactions, users will be most reluctant to accept anything less in their day-to-day activities. These trends toward exploding demand are of immediate importance to a national research and education network, but their implications are equally powerful for our common carrier network providers. A national research and education network with its test-bed and production facilities can also provide spill-over benefits that pave the way for the rapid evolution of our commercial networks.

This brief look to the immediate future helps to clarify the problems inherent in our current networking scheme. Under current conditions there are multiple "hops" before a transmission can move successfully from sender to receiver. Each hop implies that transmissions are slowed, irrespective of the bandwidth on given legs. Within a local site three or four hops are often necessary to arrive at the regional network connection. The regional then connects to the backbone and the backbone connects to the regional connector for the receiver's site. A transmission must then travel through that user's regional net to the local site. Once there, multiple hops are generally required before the transmission can reach the receiver's desktop system. Such an architecture involves inherent flaws. The number of hops is essentially fixed as long as the administration of each of the three tiers remains independent. Problems can and do arise at each hop, and it is virtually impossible for a user to know the location and/or cause of a difficulty. No central coordinating point exists through which the user can ascertain the source of trouble. No central coordinating agency exists that can seek the source of trouble, much less bring about its resolution. NSF's current, stated strategy is to provide funding for the backbone only. Though it was instrumental in bringing about the creation of regional networks, its policy now is no longer to fund them. Financial and administrative responsibility for local area nets have been and remain with the given site. With so many autonomous administrative agencies, a true system-wide solution is virtually impossible.

Yet, a system-wide solution is required. Communications must move swiftly and effectively from sender to receiver. The goal for the user is transparency of transmissions. An integrated system is required for this to be achieved. As stated at the outset, an integrated system provides such broad benefits to the society that it is fully appropriate for federal funding. This funding can be supplied to a coordinating agency through which a system-wide solution can be achieved. Without such coordination, the babble of incompatible transmissions is likely to be compounded by a rabble of battling agencies and institutions.

A National Research and Education Network: Desired Characteristics

One proposed solution would involve a greatly expanded number of nodes on a backbone plus a single administrative agency with enough "clout" to enforce standardized operating procedures sufficient to achieve significant transparency for the user. Once again, the current situation is a far cry from this. Not only is the backbone limited in the number of sites it provides, but also the regional networks are facing financial difficulties and several are faced with fragmentation resulting from state-wide networks which are just starting up. Unless a systems solution can be found, current trends suggest the potential for rapid deterioration even from current levels of service. For example, regional networks now see financial survival as their highest priority. This contrasts with a priority for service to the users. In each instance, local optimization in relation to local imperatives seems to be the order of the day. This is antithetical to the needs of the nation. Overall optimization of an effective research and education network system is an element necessary for U.S. international competitiveness.

In tuning a network the network manager can control two items, the number of hops between nodes and the bandwidth supplied at each hop. The network manager must also monitor traffic flows and provide rapid response to difficulties when they arise. At one extreme, the manager might seek connections from each node to every other node. This would minimize the number of hops, but it maximizes the number of connections and also cost. Thus within any given budgeted amount it is also likely to result in the smallest bandwidth between nodes.

In optimizing the architecture of the network, trade-offs have to be made. The objectives of various entities must be considered and history cannot be ignored. Certain principles emerge:

- The larger the portion of the total production network under the control of a single, skilled management, the greater the likelihood that an efficient, transparent network can be achieved.

- The larger the portion of the total production network that is under the control of a single, adequate financial source, the greater the likelihood that an efficient, transparent network can be achieved.

- Local and regional interests are more likely to be well served by subnets congruent with natural regional communities-of-interest, as long as those subnets offer multiple connections to the network's backbone.

In applying these principles as guidelines for a national research and education network, it would appear that they imply the desirability of a pervasive backbone with a large number of nodes, regional networks with multiple connection points to the backbone, and particular sites known as "hot nodes" directly connected to each other. Within the local sites, standards must be observed if transparency is to be achieved. Some uniformity in operating procedures will also be required. The carrot of governmental funding is likely to be the most effective incentive to achieve this. The objectives are to move the backbone closer to the strategic users, to provide greater degrees of freedom to network managers in serving such users, and to incorporate flexibility for future functionality and growth into the network.

Summary

A high-performance research and education network is an element vital to U.S. international competitiveness. However, to be effective it must be incorporated into a system-wide solution to the problems of our high-technology base. The facilities it interconnects must also be maintained at the state-of-the-art level.

The funding for the national research and education network must be consistent with the objectives of the network. The three elements, effective:

- connectivity
- transparency
- speed

require a truly integrated system. The funding approach must result in an organization that is consistent with the achievement of these goals.

The exigencies of rapid, technological evolution require that the integrated production network be coupled with an ever advancing test-bed system. The benefits for such a venture are virtually guaranteed vastly to exceed the costs.

I am happy to answer questions you may have.

APPENDIX

Government Sponsored and Other Networks Available to Cornell University

External Network Connections

Definition

External network connections provide the pathways to access mail, files, and services that exist in the outside world from Cornell and vice versa. Cornell connects to the outside world through three different types of networks: federally-funded, commercial, and user-supported. Below we describe all three of these.

Current Status

- Federally Funded Networks

The federal agencies that provide national computer networks, especially for research, are in a period of consolidation and rationalization. The agencies had chosen to pursue this goal even before the large budget cuts of this fiscal year, but the budget cuts have given them a tremendous impetus. There are now regular meetings of a Federal Research Internet Coordinating Committee, made up of representatives of the agencies which supply wide-area networking who can make decisions about budget allocations.

One of their major projects, for example, is a "Research Internet Backbone", or RIB, which will be a 45Mbps trunk across the country divided up into circuits belonging to the participating agencies. By leasing such a high-bandwidth facility together and then dividing it up, the agencies expect to save a considerable amount of money. In addition, the RIB will be available to those who need just an occasional connection as well, not just those who run large-scale computer networks.

National Science Foundation

The National Science Foundation has acquired the lead role in federal support of computer networking infrastructure for the country. Its major current project, NSFNet, is immensely popular among users and therefore overloaded. A totally new NSFNet backbone is being deployed and is expected to be in production use starting in July of 1988, using circuit speeds of 448Kbps initially and 1.5Mbps as soon as possible. In addition the NSF is the principal agency in acquiring the "RIB".

NSFNet is actually a number of networks at different levels in a hierarchy. The national backbone, NSFNet proper, is fully supported by the agency. Attached to the backbone are mid-level networks, and attached to the mid-level networks are individual campus networks. It is through the mid-level and campus networks that the NSF actually reaches its client scientists and engineers. The mid-level networks are administered separately from the NSFNet backbone, and although they have received at least partial funding from the NSF up to now, next year they will begin to become totally independent.

The NSFNet community of networks currently only support the TCP/IP protocols created for the Defense Department, although there are already plans to support the international standard ISO protocols when they are developed.

The NSF also supplies partial funding for Bitnet and CSNet, although Bitnet is almost completely a user-supported network. CSNet is supported by the NSF since it supplies "capillary infrastructure" in reaching teaching institutions and small research institutions at low cost, although it provides only mail relay and delivery services. These two networks are negotiating over possibly merging their resources.

NASA

NASA has a tremendous amount of network infrastructure, most of it devoted to mission-critical monitoring and control.

Among the NASA computer networks, SPAN is an older, although quite reliable, network, based on relatively low-speed lines (9.6Kbps, or even less, up to 56Kbps) and proprietary DECnet protocols. It is well-supported and remains popular but is no longer NASA's strategic direction.

NASNet is managed by the Numerical Algorithmic Simulation project at the NASA Ames Research Center. Its purpose is to deliver high-speed communications to collaborators on the NAS project, and it consists of mostly 1.5Mbps trunks connecting together ethernet segments.

The NASA Science Internet (NSI) is the youngest of the NASA networks of interest to Cornell. NSI is still being deployed, and is strongly affected by the current mood of resource sharing in Washington. The current model is that NSI will supply high-speed (at least 1.5Mbps) trunk circuits, but in many cases will depend on the NSFNet mid-level networks to reach the actual researchers. NSI currently supports only the TCP/IP protocol suite. The NSI project office is also at the Ames Research Center.

A major gateway between SPAN, HEPnet/PHYSnet, and NSI is planned for Marshall Spaceflight Center by the end of 1988.

DARPA

The Defense Advanced Research Projects Agency runs ARPANet, the original and still a pervasive national research internet. Because of recent budget cuts, and because ARPANet is based on old technology (the first node was delivered in 1968), in recent months support for ARPANet has been seriously reduced. Circuits are being removed, nodes which have other access to the TCP/IP Internet (especially through NSFNet) are being disconnected, and facilities are being merged with the Wideband Net, another DARPA project.

A new network is being developed at DARPA: The Defense Research Internet (DRI). This network will support research into high-speed networking, the original intent of ARPANet. The NSF is emerging as the prime agency in supplying network infrastructure, but DARPA is continuing in its role as the lead agency for research in network technologies.

DoE

The Department of Energy runs MFEnet and HEPnet. MFEnet is for researchers into fusion energy. HEPnet, which uses DECnet protocols and is transparently connected with SPAN, is open to those engaged in high energy physics research. MFEnet is based on relatively old technology, and the DoE is in the process of developing a new network called ESnet (for Energy Science Network) which will be able to carry protocols from both of the older networks, plus TCP/IP and the emerging ISO protocols.

• Commercial Network Services

Unlike the networks described above, commercial networks such as Tymnet, Telenet, and Accunet primarily supply low-level X.25 communications support. The users are then free to build upon this base, but the only actual application-level service supplied is terminal support. Some value-added functions such as IBM 3270 emulation are available. Since they are public facilities throughput is usually limited although Accunet claims very low response time between directly connected nodes. Users pay a basic rate as well as per character or packet of characters transmitted.

Carriers are currently experimenting with offering full-service networks. NYSERNet, the New York State Education and Research Network, was developed with partial support from NYNEX so that NYNEX could gain such experience.

There are other facilities, such as switched 56Kbps service, which Cornell has never used.

• User-Supported Networks

Usenet, originally developed for UNIX sites, supports mail and ASCII file transfer over dial-up telephone connections. Addressing on the net is pass-down-the-line, with each intermediate host explicitly specified in the address. Thus, to send mail, you must know the exact path to the destination machine. The two main functions of Usenet are mail and Netnews. Netnews is an international bulletin board system, supporting hundreds of interest groups (newsgroups) with topics ranging from UNIX support to Macintosh software to movies to cooking. The network has extended well beyond its original "UNIX-only" nature, and is supported by a wide range of machines and operating systems. Mail transfer now takes place over the Internet as well as dial-up phone lines. The only requirement for joining the network is finding an existing member willing to serve as your link. There is no central organization at all although volunteers coordinate among themselves.

Although Bitnet has a small amount of NSF support, it is almost entirely supported by its members, degree-granting institutions and their research partners. The only requirements for joining Bitnet are that you be engaged in research, lease a communications line to an existing member, and be willing to allow at least one other new member to connect through you. There is a membership fee. Bitnet uses proprietary protocols and thus is not well-connected to the other national networks except by mail relays. Bitnet offers only file transfer and interactive real-time messages; it is currently experimenting with establishing a TCP/IP backbone as well as with running its protocols on top of TCP/IP using the NSFNet backbone. This would allow small Bitnet sites full TCP/IP Internet access in addition to their current proprietary protocol Bitnet access.

- **Cornell's Connections to the National Networks**

Cornell is a major node on the NSFNet backbone for several reasons. First it is the home of the Theory Center and the National Supercomputing Facility, one of the six NSF-funded supercomputer centers. Second, it is the connection point between NYSERNet (the New York State Education and Research Network, an NSF mid-level network) and the NSFNet backbone. It will also soon be the connection point between NSFNet and the Ontario TCP/IP network, and probably a major gateway for NRCNet, the proposed Canadian backbone net. Finally the Theory Center networking division is a major contributor to the NSFNet backbone project, supplying software, engineering and general consulting. We are well-connected, in many ways, to NSFNet.

We are also a major node on Bitnet, and are participating in the Bitnet TCP/IP experiments. All of Canada is currently connected to Bitnet through Cornell.

Because of its connection as a major node of NSFNet, Cornell has given up its CSNet membership. As mentioned, CSNet exists mostly to reach lower-budget institutions at low cost; Cornell has all of the functions CSNet could supply with better performance through other means.

The Spacecraft Planetary Imaging Facility at Cornell has one host attached to SPAN via a 9.6Kbps circuit, but at the moment no other host at Cornell can use this link because SPIF only has enough money to support an *end-node* connection.

We currently have a 56Kbps connection to Arpanet through the Theory Center and a 9.6Kbps connection through Computer Science (which we are giving up). Many sites are losing their ARPANet connections, but because we have a supercomputer center here we are guaranteed a connection as long as it allows NSF researchers to reach us who could not reach us by any other network. We do not have plans to be part of the Defense Research Internet at this time although there is a significant chance we will do so.

Wilson Laboratory is part of HEPnet (currently only at 9.6Kbps), and there is a definite chance that all of Cornell could access HEPnet and SPAN through them. This possibility can be explored as needed.

An MFEnet connection comes into Clark Hall but, similarly, cannot be used as a gateway for the rest of the campus.

Cornell currently has connections to Tymnet and Telenet commercial network services. We use these only for access to the IBM systems and some limited access to the Theory Center's non-IBM systems. These networks are useful to us because they allow terminal-style access from areas we would have difficulty reaching otherwise, such as Africa, but the federally-funded research networks will have more useful functionality for the foreseeable future.

A number of machines at Cornell are on Usenet, and Netnews is widely available. Many members of the Cornell community participate in Netnews discussions. In addition, Cornell acts as a Usenet link for several local businesses.

In summary, Cornell is quite well-connected to the national networks, and is in fact highly involved in helping to create a better national network infrastructure.

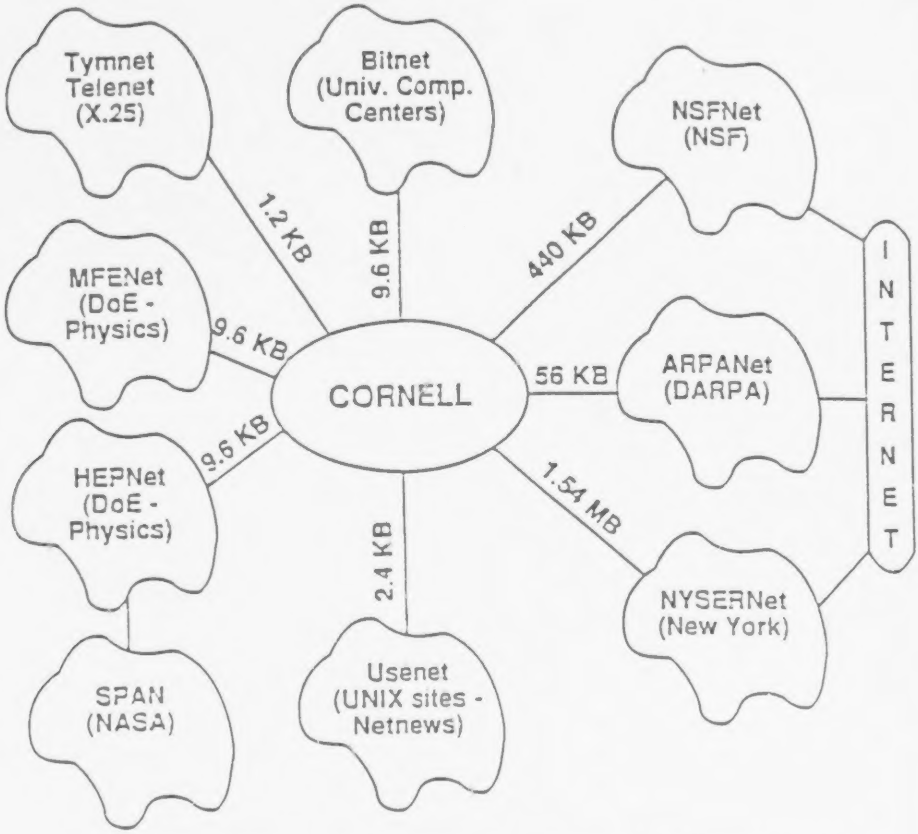
Tymnet
Telenet
(X.25)
Bitnet
(Univ. Comp. Centers)
NSFNet
(NSF)
MFENet
(DoE - Physics)
HEPNet
(DoE - Physics)
SPAN
(NASA)
Usenet
(UNIX sites - Netnews)
ARPANet
(DARPA)
NYSERNet
(New York)
CORNELL
INTERNET
1.2 KB
9.6 KB
440 KB
9.6 KB
56 KB
9.6 KB
2.4 KB
1.54 MB

Current Status of Cornell's External Connections

QUESTIONS OF THE CHAIRMAN AND THE ANSWERS THERETO

Q-1: Isn't there an inherent unfairness to a private network supplier from having a federally supported network providing services to the research and education community?

A: The short answer is no.

Properly implemented, the two types of networks would not be competitive with each other. In economic terms the National Research and Education Network is a complement to the commercial networks, not a substitute for them. First, most likely it would be implemented through the resources of the commercial networks themselves. Second, it's users are unlikely to represent significant revenue potential for the commercial networks. Third, a federally supported network will create, identify and aggregate new categories of users for the services of the commercial networks. It will also help to advance the state of the art of networking in general. These points require some elaboration.

In the first report of the Market and Economic Impact Study Group (MEIS) which I headed for NYSERNet, we analyze this problem in the context of NYSERNet, a regional education and research network, and its interaction with New York Telephone (NYTEL), the regional commercial network. The analysis has general applicability and can be extended to a National Research and Education Network.

The first element of complementarity is that NYSERNet was implemented over lines leased from NYTEL. Next, in our study we identified three categories of users of telecommunications services and analyzed the relationships among them. The first category is made up users with characteristics similar to those of likely users of the national network. They offer little commercial potential. The other two categories are made up of very likely new customers for the commercial networks The following quotation from the executive summary of that report [with added explanatory notes] lays out that analysis:

> "We have identified three categories of NYSERNet users:
>
> Category #1: Users at NYSERNet node-institutions who have access to the full high-speed capabilities of NYSERNet through their institution. [Note: These users are mainly research oriented.]
>
> Category #2: Users with access to the functions of NYSERNet through a Centrex LAN [data-over-voice @ 19.2K baud] connection sponsored by an institution that is not a NYSERNet node. [e.g., a local school district. These are new users for these services of the commercial networks.]
>
> Category #3: Users with access to the functions of NYSERNet through a Centrex LAN connection which is obtained by them as individuals. [These also are new users for these services of the commercial networks.]
>
> "Our research established that Category 1 users:

> - Represent the main test bed for new products and services that can improve network performance. These users are in a position to press their needs for services that are easier to use and facilities that are more reliable. They are the users most likely to bring about long-term network improvements that will benefit all users. They therefore represent longer-term revenue potential.
>
> - Face many problems in carrying out computing tasks using networks. Until these problems have been overcome, the numbers of such users will remain limited.
>
> - Represent relatively small revenue potential for NYSERNet [or a commercial carrier] in the short-run.
>
> "By contrast, our research established that Category 2 and 3 users:
>
> - Face relatively small problems of interconnection and linkage to NYSERNet via Centrex LANS. They use mature technology.
>
> - Represent relatively large revenue potential for NYSERNet [and the commercial carriers]. in the short-run.
>
> "We conclude that persons from Category 1 institutions, in their roles as individuals, constitute a large, likely, potentially fully compensatory market for NYSERNet services [and the commercial carriers] -- as Category 3 users. In other words, Category 1 institutions geographically aggregate qualified Category 3 users who, in the short run, offer a potential market broad enough to sustain financially the initial phases of the development of the network. We conclude further that there is an additional group of likely users associated with other, non-node institutions who also represent significant short-run revenue potential -- Category 2 users.
>
> "These facts lead directly to our action recommendations, which make maximum use of resources which NYNEX and NYTEL already have in place or can rapidly bring to bear (e.g., underutilized Centrex facilities in parts of Manhattan and other areas of New York State)."

The Category 1 users--those who would be served by the National Research and Education Network--represent only insignificant revenue potential for the commercial networks in the short run. Nonetheless the commercial networks benefit from their existence since the National Research and Education Network would be implemented through the facilities of the commercial networks that would be paid for by the federal government. The Category 2 and 3 users represent a substantial new market for the regional service companies. These customers would not exist in the absence of the National Research and Education Network. Their need for service derives from their association with the main nodes on the Governmentally supported network. The interactions of the two types of network are complementary, not competitive; they reinforce each other; they do not compete for the same categories of users.

This relationship is consistent with the economic characteristics of the services being provided. There are huge societal spill-over benefits from the research activities of category 1 users at the main nodes on the network. This is the activity appropriate for government support, but its implementation will be through facilities supplied by the commercial networks. The use of the network for primary and secondary education by the Category 2 users is fully appropriate for state funding, and implementation will again be through the commercial plant of the private phone companies. For Category 3 users the benefits of the connection from the home to the main node is one of personal convenience as well as productivity. The economic benefits of the activity are essentially capturable to the individual party and to the individual's particular research project. Under these conditions normal market incentives should apply--i.e., fee-for-service pricing through the commercial networks.

Q-2: What about the analogy to the National Highway Network? Is there a way to fund the National Research and Education Network through "user charges"?

A: Perhaps, and the rationale grows from an analysis similar to that in question 1 above: institute a "user-fee" on complementary products that are only indirectly related to the National network; use those revenues to fund the National Education and Research Network that serves users like those in category 1 of our study. We recommend Federal payment from general revenues, however. To see why we make these recommendations, let us develop the analysis in part by exploring the analogy between a National Research and Education Network and the National Highway Network.

Despite the "user charge" implemented to fund the National Highway Network, that network of highways is perceived by users as available to them "free". This is because the user-fee comes in the form of a tax on a complementary product, gasoline. The user does not make a direct connection between the user-fee and use of the Highway Network since the fee is not direct. The user-fee applies each time gasoline is purchased, rather than being limited only to those times a highway in the national network is used. Also the fee applies to all users of gasoline, even those who do not drive on the interstates at all.

As suggested in the answer to question 1, the closest complementary products (services) related to the National Research and Education Network are those used by Category 2 and 3 users. If we push the analogy with the highway network, a user-fee applied to these users would be a way to achieve funding for Federal Government support of the National Research and Education Network. Given that there are existing excise taxes on telecommunication services, a portion of such a "user-fee" is already in place. It would be automatic with increased telecommunications use and might only require augmentation. To make the analogy closer, yet, the new revenues from the excise tax on these particular users could be earmarked as a "trust fund" to be applied only to the funding of the National Research and Education Network.

As we argue in our paper, "Economic Benefits and Public Support of a National Research and Education Network", however, such a national network is fully appropriate for direct federal support from general revenues. The rationale is repeated in the paragraphs below. A middle ground which represents the closest analogy to what is done in the case of the interstate highways would be to create a trust fund based on an excise tax on all telecommunications services as the means for raising the revenues necessary to support the National Research and Education Network.

These comments are based on fundamental economic principles: there is a spectrum of economic goods and services with different susceptibility to being efficiently allocated through private, competitive markets. Two criteria apply. First, who gets the benefits from the product as compared to who pays the costs for it and, second, how do you achieve efficient allocation of the good or service once it has been produced? If all the benefits accrue to the party that pays for and uses the good or service, then it can be allocated efficiently through competitively structured private markets. Such goods are called "economic private goods." But if the benefits from a product or service (e.g., knowledge about high-temperature superconductors once that knowledge has been produced) accrue broadly and automatically to the society as a whole and at the cost only of reading the articles reporting the breakthroughs, while a private party (e.g. IBM Europe or Rice University) must incur the cost for its creation, a market can not efficiently allocate that product or service. Such goods are called "economic public goods". The potential benefits from such goods are great and broadly dispersed. Thus it is highly desirable to bring them into being. But they can be efficiently allocated only if the price for their use (use of the knowledge about superconductors) is equal to the cost (in economic terms, the Marginal Costs [MC] of making them available to additional users). The solution to the apparent paradox of (high) costs of creation of a good or service (knowledge about superconductors) and the need for a price close to zero for its use, comes through a broad-based source of revenue for society as a whole (a tax not directly related to the product or service itself) made available to the producer (through a research grant, for example), while making the product (knowledge about superconductors) widely available and essentially free to users (through publications read by additional research persons).

This appears to be (and to have been) the situation for the Highway Network which has huge spill-over effects for the society as a whole and, once built, small MC for each new instance of use by an automobile. The approach to funding the national highway network has been chosen to be only indirectly related to the use of these particular highways such that they are seen to be free. The National Research and Education Network represents the same pole in the spectrum of economic goods and services as do these highways. By definition, the Network is used to facilitate activities--research and education--with enormous spill-over benefits that can not be captured directly by the party who must pay to create them. Here also the market fails as an efficient allocator of resources. This network should be dealt with through "asymmetric pricing", as the highway network has been: the price to the user is zero (it's free), while the large costs necessary to bring the highways into being are covered through a broad, general tax that is not directly related to the use of the highways (the gasoline tax). For the National Research and Education Network a similar approach should be used, ideally with a tax based on general revenues. And now a brief digression.

Historically in the case of the private telecommunications network, another type of market failure has also existed: that known in economics as a "natural monopoly". It occurs between the polar cases of economic private goods and economic public goods. This phenomenon occurs when there are great scale-economies in providing services, such as telecommunications services, through a single facility (e.g., the network). The result is that the average cost of supplying service to every user, while not zero, falls as additional users are added (to the network). Under such conditions, it would not be economically efficient to create a competitively structured market. Two or more networks would impose higher costs on everyone. That is why a monopoly franchise originally was given to the phone company in exchange for public utility status and regulation. The regulator's role was and is to assure "equity" and "reasonableness" in the pricing done by the monopolist, while assuring the lowest cost way of providing the service for the society. Ideally the system provides service at the minimum average cost to society with only a "normal" return being earned by the private monopolist--as a result of the constraints imposed by the regulators.

Increasingly, with changing technologies in the more recent past, it is possible for economic competition to take place among suppliers of longer distance telecommunication services. This is an industry, then, which provides the full spectrum of "economic goods":

- private goods, those that can be efficiently provided through competitively structure markets: e.g., the telephone equipment itself, and increasingly long distance services;
- natural monopoly services, those that can be provided most efficiently through a regulated monopoly: e.g., the "last mile" local connection to the users' premises;

- economic public goods, those with an extremely high ratio of spill-over benefits to capturable benefits: e.g., a Research and Education Network.

If a telecommunication service has a high proportion of benefits capturable directly to the user (and, we assume, "natural monopoly" conditions are not present), then the supplier will be able efficiently to supply this service through markets in competition with other suppliers. To the extent that services have this economic characteristic, they can be supplied efficiently by private firms operating in a competitive structure. In this country we provide natural monopoly-type services through private, regulated, monopolies. For economic public goods, private markets of any kind are very inefficient. That is, when the spill-over benefits greatly exceed the privately capturable benefits, producers of the good/or service are unwilling to pay for more than that portion of the total benefits that they are able privately to capture. Despite the existence of enormous spill-over benefits, the product would not be offered beyond that which can be sold to the private parties experiencing the private portion of the benefits. This is true, unless of course, an agent for the society as a whole provides sufficient resources to the supplier, to bring the full benefits into being.

Within the spectrum, the National Research and Education Network is, as stated above, at a pole. It is neither efficient nor appropriate to attempt to provide it through a market or a market-like structure. This is a textbook example of the type of service that should be provided through asymmetric pricing--payments to the supplier from broadly based tax revenues not directly related to the use of the service, and a price of zero to the user. The society as a whole benefits from its existence, the society as a whole should provide the resources through general tax revenues to bring it into being. The broader the base of the tax, the better. An excise tax on telecommunications usage to fund the National Research and Education Network--with the tax earmarked through a trust fund to be used for the network--would be analogous to the gasoline tax for the highway trust fund and an acceptable approach. We still recommend general revenues for funding the network, however.

Summary

Let us now pull together the answers to questions (1) and (2). In the response to the first question, an analysis was presented demonstrating the interrelationships among the levels of service provided through a network such as a National Research and Education Network. It was pointed out that where the use of the network is in the furtherance of research and development activities--themselves activities with the characteristics of economic public goods--no market can efficiently allocate the fundamental "product" produced through use of the network. Nor is there a market for the services of the network itself: the network creates enormous spill-over benefits to users whose products--research and education--create enormous spill-over benefits to the society as a whole. Normal market forces are not present. No attempt should be made to impose them. The user can not "sell" enough of his or her product at this stage in the innovation process to sustain the network. Asymmetric pricing is economically efficient: zero price to the user, payment to the supplier from broad-based tax revenues.

The clearest case at the other pole, economic private goods, is that of the Category 3 user using services whose benefits--flexibility in work patterns and work location--accrue in greatest measure to themselves and to their immediate employers. The infrastructure for billing already exists, fee-for-service is appropriate here and should be used. To the extent that there are some spill-over benefits that accrue to a given research or education project through improved productivity of the worker, the project itself could find it "profitable" partially to cover the costs of network connection to the worker's home.

The implication, then, is: the closer the economic characteristics of the service provided by the network supplier to being a "economic private good", the more appropriate the use of market mechanisms through the commercial network for allocation of the service. The ground between the polar extremes in the telecommunication industry represents a continuum which is currently in a state of transition and flux. An increasing portion of the services provided by telecommunications networks is capable of being supplied through competitively structured markets. The portion which is the subject of this hearing is not. It represents a polar extreme. The dividing line between that which is appropriate to be provided through federal government support and that which is not, is determined by the fundamental economic characteristics of the service being provided. Those activities which are susceptible to being provided through normal market mechanisms do not require federal support. Those with disproportionate spill-over benefits do.

Q-3: In order to use the National Research Network to enhance our competitiveness, we'll want to have industry linked to the network.

- Do you think that businesses will see the benefits of belonging to a National Research Network?

- Do you think they will be willing to pay for the privilege?

- Which businesses do you think will be first to join?

A: First, short answers:

- Yes, businesses will see the benefits of belonging to a National Research and Education Network.

- To some extent they are likely to be willing to pay to connect to the network, but to the extent that the activities that they engage in also have a very high ratio of spill-over benefits to privately capturable benefits, it would be inappropriate for them to do so.

- A number of businesses have already joined the prototypes to the National Research and Education Network. From this and other information the most likely classes of participants can be identified

The analytic underpinnings of these short answers are presented in the responses to questions 1 and 2.

Also contained in the prior responses is a basis for questioning the assumption implicit in the statement of this question, question 3. The National Research and Education Network should provide significant benefits to U.S. international competitiveness even if businesses are not connected to the network. To the extent that the National Research and Education Network enhances the productivity of both research and education in our university and laboratory systems, this can significantly enhance the international competitiveness of American business. To the extent that by connecting to the network itself, research and education within the business firms of the country is further enhanced, this would be an additional benefit to U.S. international competitiveness.

Explicit in the responses to questions 1 and 2 is the analytic framework of a continuum of economic goods and services from the pole of purely private goods--all of whose benefits and costs are privately capturable to the participants in the purchase/sale transaction--to the alternate pole of purely public goods--whose benefits become broadly available to all parties and virtually free whether or not those parties have been associated with the costs of creating the public good. In other words, there is a continuum of goods and services from those with no spill-over benefits to those all of whose benefits are essentially spill-over benefits. Basic research produces mainly spill-over benefits. Applied research tends to produce benefits, some portion of which is capturable to its creators, but with continuing, significant spill-over benefits. Thus for basic research activities, it would be appropriate for business firms to participate fully in a National Research and Education Network that is Federally funded. For their applied research activities, business firms should be willing to pay--and should be charged--fees which are consistent with the degree of private benefit the firms are anticipated to receive. Conceptually, that portion of the benefits that represents spill-over benefits should continue to be subsidized.

This analytic framework provides guidelines for identifying the portion of total activities which should be subject to fees directly paid by users. It leaves open the extremely difficult tasks of identifying each type of benefit and implementing a system consistent with capturing fees only for the privately capturable components of benefit.

The types of businesses that have been the first to join the prototypes to the National Research and Education Network have been:

(a) those that produce products used in networking services,
(b) those that have a commitment to basic research.

Firms in group (a) see the network itself as a living laboratory through which to carry on research and development for the products that they produce. This is a symbiotic relationship. The National Research and Education Network can serve as a testbed for the creation and enhancement of computer and networking products. Once proven and demonstrated through the testbed, these products can be sold to private users as largely private goods. Firms in group (b), to the extent they are engaging in basic research, should be full fledged participants in the National Research and Education Network. To the extent that they are engaging in applied research, some of whose benefits are directly capturable to them, they should pay an appropriate portion of the costs of the network.

The comments thus far suggest the desirability of establishing the Research and Education Network as a limited access network. It should be available to practitioners of research and education whose outputs embody a high proportion of spill-over benefits. It should not be used for the day-to-day services in support of the private goods produced by business firms. Activities of the latter sort should be carried on through commercial telecommunications networks on a fee-for-service basis. Guidelines for use of the National Research and Education Network must be prepared that establish the boundary with a threshold for including those activities with significant, but not dominant spill-over benefits.

One of the major payoffs for a National Research and Education Network is the rapid, almost instantaneous availability of new information and knowledge as it is created. Our universities and research laboratories are institutions created to provide output of this kind. The motto, "publish or perish," embodies the ethic for these institutions. Those who "give away" the best ideas soonest, are those most highly rewarded. This ethic is in direct conflict with the requirements for successful business (or national security) operation. Businesses seek new, innovative products whose benefits can be captured privately for the enrichment of their stockholders. Applied research and development activities result in outputs, some of whose characteristics are appropriate to each camp. This presents a dilemma for effective, efficient interaction between the two groups. The same dilemmas apply to network connection and access to the information and knowledge created in and by each of these types of institution. For activities that fall in the gray area between them, institutional mechanisms must be developed which honor the ethic appropriate to each.

Q-4: How are we going to pay for this new network? There are estimates that by 1995 researchers will be spending up to $1 billion on networking. That estimate assumes that there will be one or two million researchers and students using the network. The total networking bill could be comparable to the amount that today's researchers spend on telephone calls.

- Do you think these billion-dollar estimates realistic?
- Should the federal government pay the whole bill?
- Do we need to charge individual each user of the network to keep it from being overwhelmed?

A: Again short answers:

- The $1 billion estimate is likely to be low in relation to the full costs incurred by users of the National Research and Education Network. It may be reasonable for that portion of the total costs for which it is appropriate for the Federal Government to pay.
- As laid out in the analysis above, the Federal Government should pay the full costs of the network whose use represents "economic public goods". Even here, however, the costs which are perceived to be "the cost of the network" are likely to include only the wide-area telecommunication network itself. Users will incur costs for their desk-top devices and local-area networks independent of the wide-area network. These are likely to represent multiples of billions of dollars worth of equipment.
- Two actions should be coordinated to keep the National Research and Education Network from becoming "overwhelmed". First, the use of the network should be limited to those activities with a very high proportion of spill-over benefits. Second, the capacity of the network should be maintained at a level sufficient to provide service commensurate with these spill-over benefits which accrue to the society as a whole.

These short answers are elaborated below:

A very significant infrastructure already exists within our universities and research laboratories which will permit access to the National Research and Education Network. Additional resources will be committed to upgrade these existing capabilities as the National

Research and Education Network comes on line and becomes enhanced. In reality, all these resources constitute true components of the network--after all, it does no good to move information and data from one point within a network to another point. Rather, information and data must be moved from one user to another; their connections to the network are necessary components of the network.

The discussion of the portion of the bill that should be paid for by the federal government and the rationale for determination of that portion have already been presented: to the extent that there are societal benefits which would not be brought into being except through funding from the Federal Government, such funding is fully appropriate.

Next comes the question of allocating the resource--the network--after it has been built. First, the National Research and Education network itself has characteristics of an economic public good. Once the network is in place its use by an additional user involves costs only so negligible as to be considered "zero". Economic efficiency requires that the price charged to the user be equal to the marginal cost the user imposes on the society (i.e., on the network). In this instance the price to the user should be zero--use should be "free".

There is an additional practical reason for providing the services of the National Research and Education Network in a manner perceived by users as "free". There are enormous costs to the identification of use of the network to a given party and the billing of that use to that party. Significant human and technical resources must be devoted to devising algorithms for billing such use to the individual parties. Current estimates are that billing costs for a commercial telecommunications network exceed the magnitude of all other costs combined. That is, the effort required to bill customers on a fee-for-use basis more than doubles the costs of providing the service. Neither the hardware nor the software for identifying to a given user that party's use of a packet switching network has yet been developed. These facts are important to the rationale for the third short answer given above.

It is more efficient to follow simple administrative practices for limiting the use of the National Education and Research Network to those classes of users that provide spill-over benefits to society, than it is to attempt to charge users on a fee-for-service basis for each use that they make. The resources saved by avoiding the accounting and billing process can much more efficiently be devoted to providing additional capacity to the National Research and Education network as appropriate, authorized use grows. That is by far the better way to prevent the network's resources from "becoming overwhelmed".

In the answers to the earlier questions we have developed the analysis that leads to the conclusion that a "user fee"--instituted on indirectly-complementary products: telecommunications use in general--is an appropriate approach to raising the revenues necessary for the Federal Government to pay up front for the Research and Education Network, but our preferred approach remains: pay for the Network from general revenues.

Senator GORE. Thank you very much. I am going to ask only a few questions of this panel so we can all get some lunch and end the hearing. I do want to express my appreciation to all three of you.

Dr. Connolly, when you were at NSF, you helped to start NSFNET and since you left, the work has continued, the capacity has increased thirty-fold. Now that you are back doing research, you can give us some idea of how well the network is working from your perspective.

You have done that in your statement, but could you give us more details? Is NSFNET providing you and your colleagues at the University of Kentucky with the supercomputer access that you really need to do your research?

Dr. CONNOLLY. Yes. I use it myself all the time. Many of my colleagues are plugged into our own supercomputer through the local network but then they are also for various other purposes plugged into the Pittsburgh center, the Illinois center. The system that has been set up by NSF is a very user-friendly one.

The difficulty comes when it becomes very, very popular. The word gets around that this is a wonderful thing to have for a research tool and then pretty soon it is clogged up in the middle of the day and you cannot get logged on.

Senator GORE. That has already happened?

Dr. CONNOLLY. Well, it happened. It got very bad in June and then suddenly got better in July because of the upgrade of the network. We saw a marked improvement and service at the moment is very, very good, but we know it is going to fill up again.

One of the problems is, as I pointed this out in my testimony, is that there is no control over network access. Now, if you suggest a user charge, I tend to disagree with that. I think a better way is a controlled access through allocations of network time. That is what we do with supercomputers.

Every supercomputer center gives time away free but they do not give infinite amounts of time, 10 hours, 100 hours, on the basis of a peer review process. We do not do that with networks. We give a person a network address. If they use 100 hours a week or if they use ten minutes a week, we do not even know. The controls are not there in place. They should be. They could be. The technology is possible.

Senator GORE. Since it is Federal dollars for the most part, it is an allocation problem and I am certainly open to different ideas on how to solve that problem.

If you have this enormous new capacity and you look at a hybrid research and commercial use, then one is attracted to some kind of user fee as a way to have a common allocation scheme for both kinds of users. That is sort of the reason I was headed in that direction. But as I say, I am open on the question.

On another point you referred to the need for a lead agency and in your written statement you say that ideally NSF should be the lead agency.

And yet NSF's mission is to fund basic science and not necessarily to run research facilities. And, again, I am open on the question but are there alternatives you see to making NSF a lead agency for network?

Dr. CONNOLLY. Let me just point out that there are certain branches of science where research facilities are absolutely essential. A good example is astronomy. NSF funds several astronomical research facilities. Other examples are cyclotrons and particle accelerators. There are many, many branches of modern science where big machines, big facilities are absolutely essential.

You just cannot make progress in these areas without facilities. Computational science is similar. You have to have a normal computational facility and I, myself, do not really distinguish between the centers and the network. Having the centers without the network is like having a car without wheels. And a network without the centers is vice versa.

They are all a part of the whole. It is a national research facility and I think it is very appropriate for NSF to run it. Besides, nobody else wanted to do it.

Senator GORE. Dr. Wulf, you certainly agree with that, do you not?

Dr. WULF. First, let me point out that NSF does, in fact, support a number of research facilities, including telescopes and ships, for example. Perhaps the major difference here is that the supercomputer centers and the network are applicable to all of science, whereas the other facilities tend to be single discipline in nature.

Second, I do not believe that NSF should be a "lead agency" in the sense of getting all the funds and having total responsibility. NSF is serving the role of *de facto* lead agency as first among equals with respect to deploying precommercial, but nevertheless available, technology. That does seem appropriate; somebody had to step up and do that job.

DARPA, by contrast, has assumed the role of *de facto* lead agency with respect to advanced networking research. They have a long history there and I think that is also totally appropriate. The various agencies have different needs and specialized user communities. It would be counter-productive for NSF to take over their special networks.

The key is interoperability between the various component networks, and we are working in that direction. Essential to that interoperability is getting a common set of protocols, for example, and we are almost there.

Senator GORE. Well, again, I have an open mind on it and as Dr. Connolly said, nobody else would do it so NSF stepped in. I certainly am grateful for what has been done so far. At the same time, I heard Dr. King's strong statement that you need somebody who is actually going to do this. I think all of us need to take a good close look at where the capacity is and what the institutional constraints are.

I look forward to sharing a dialogue with you and others on that in the time ahead. Let me move to another question for you.

When the OSTP report came out last November, it was too late to have much effect on the fiscal year 1989 budget cycle. However, the head of OSTP, Bill Graham, said some of the proposals would receive funding through osmosis. In other words, the moneys would diffuse out of other parts of agency budgets and into computer research and development. Needless to say, that created some alarm on the part of scientists in other areas struggling with very tight budgets.

How much osmotic funding do you anticipate there will be at NSF in fiscal year 1989?

Dr. **Wulf**. Maybe I could get back to you with a more precise answer to that.

As I said, we are increasing the networking division budget by about 50 percent in 1989; that, I suppose, could be considered to be partly osmotic funding. However, it probably more accurately represents the fact that the Foundation, and my predecessor, recognized the importance of networking and really anticipated the study.

Senator **Gore**. Well, NSF representatives were on the FCCSET Committee that wrote the report, so I suspect that NSF is generally pleased with the report, but are there any parts of the report that the foundation is not happy with?

Dr. **Wulf**. We are very happy with the report. We are not always pleased with the myopic reception of the report, which tends to focus on its supercomputer and networking aspects and forget that the other two legs of the stool are equally important.

Senator **Gore**. Dr. McAdams, as part of your work you have examined the economic impact of many of the new developments in computer technology and you have focused on economic benefits of the networking in your testimony.

In your opinion, has networking been getting an adequate slice of the funding pie?

Dr. **McAdams**. Absolutely not.

Senator **Gore**. Can you be more succinct than that?

Your answer should be longer. I do not want you to run on and on. The National Academy of Sciences report indicated that most of the obstacles to developing a National Research Network are managerial and you have discussed some of the problems with the present network.

If we are going to have one organization responsible for coordinating and implementing a national network, who should it be? Should it be NSF, should it be a new entity created just to run the network? What is your opinion on that point?

Dr. **McAdams**. My feeling is that——

Senator **Gore**. You have already made Dr. Wulf mad earlier with your comment about NYSERNET, so do not worry.

Dr. **McAdams**. Just local pride, local pride. I think we need to have an overall systems approach to our technology base. I think this network is one piece of it. I think we need a civilian technology agency that provides overall coordination for civilian technology. I do not think we can do this halfway.

I have tried to stress the problems that are inherent in doing it halfway. Any time you have a system problem and you change only one element of the system, you can be virtually assured of lack of success of your efforts.

If we make a wonderful network, I am sure the Japanese will be more effective in taking the R&D out of that network and implementing it into their products than we are.

Senator **Gore**. Whether it is the electronic gateway on the fiber optic cable or whether it is printed reports that come out.

Dr. **McAdams**. They will find a way.

Senator **Gore**. The first big component of this larger task is applications, correct, the step between R&D and the commercialization. That is where we are falling down. This is another whole hearing.

Dr. **McAdams**. Right, I agree, it is another hearing but it is hard to respond to the question without putting the answer in context.

Senator **Gore**. Surely you are not saying that the entity in charge of running the network ought to be one that is also in charge of answering this applications question?

Dr. **McAdams**. I put it the other way. If we had the overall entity within that entity and under its control, there should be then a subgroup that is organizing the research and education network.

Now, I think it is important to point out that we do have a three-tier network now and the three tiers will not work into the future, that when we are making rapid transfer of graphics and imaging and so forth, we cannot have a back one which is managed by one backbone group, regional networks, each managed separately by the regional groups and local networks under the control and operation and differential funding management capability and support at each of the local areas.

That is what I meant by the problems of having end to end communication.

This has got to be a system. It has got to be a network system, not a series of these three tiers.

Now, with NSF budget limitations, I do not think they had any choice. Thus, my answer to your first question is absolutely insufficient funding for this networking approach. When NSF is putting a region afloat for next year to do all their own funding, I think that we have a very substantial problem. We certainly have a problem of focus because the regionals are now going to have to focus on surviving and they are surviving in an economic market where fundamentally they are disabled. There is no market that works there.

Senator **Gore**. Okay. Well, I want to thank everyone. We are going to have to call it a day. We have a lot of work to do on this Subcommittee and all of the institutional representatives here have made us aware of the agenda that needs to be addressed by them and by us. I just want to pledge to all who have taken the time and showed the patience to be here all day that as one individual Senator and as a member of this Committee, I am going to be continuing to pay an awful lot of attention to this and devote a lot of time and effort to it. And I welcome any suggestions from any of you in the weeks and months ahead as we plan for the next year's budgets and authorization bills and any new initiatives that need to be undertaken. I am available as a resource and anxious to help, and I appreciate all the work that is under way.

Thanks to everyone.

[Whereupon, at 1 p.m., the subcommittee was adjourned.]

ADDITIONAL ARTICLES, LETTERS, AND STATEMENTS

STATEMENT OF DR. THOMAS P. RONA, DEPUTY DIRECTOR, OFFICE OF SCIENCE AND TECHNOLOGY POLICY

OSTP is pleased to offer the following for the record:

OSTP, through the FCCSET (Federal Coordinating Council on Science and Engineering Technology) Committee on Computer Research and Applications, published "A Research and Development Strategy for High Performance Computing". This report was requested by P.L. 99-383. OSTP organized the FCCSET Committee into three subcommittees:

Science and Engineering Computing; Computer Research and Development; and Computer Networking, Infrastructure and Digital Communications.

OSTP submitted the report on November 20, 1987. Public Law 99-383 originally called for two separate studies: Year 1 was to study networking of supercomputers—computers used in research; Year 2 was to study networking of all research computers. The FCCSET Committee felt that to single out networking to the exclusion of other computing research areas would potentially lead to distorting of research priorities. It was agreed to expand the scope of the report to include the relevant context of research computing; and simultaneously the time frame was reduced from the original 2 years. The OSTP report therefore places networking needs within the overall context of the research infrastructure. The report identifies the need for investment balance among high performance computing architectures, software to enable their productive use, networking to ensure broader research access, and needs to enhance the research infrastructure. The National Academy of Sciences subsequently reviewed the networking element and published "Toward A National Research Network" in July 1988.

The FCCSET is currently adressing itself to the Congress's request to effect a quantitative assessment of the benefits resulting from implementing this strategy. This response will be in two phases: the first, shorter term, more anecdotal to accompany agency budgets; the second, longer term and more quantitative. Another part of the Congressional request related to the quantitative requirements for high performance computing networks has been completed by submitting the backup documentation of the OSTP Report.

Agencies are currently involved in developing their budgets; cross-agency coordination will be done within the FCCSET. OSTP is expected to review this by October 1988. The Director, OSTP sent a letter to agency heads indicating OSTP's support for agency budgets in the high performance computing area within the overall context of the development of the President's budget.

Toward A National Research Network

National Research Network Review Committee
Computer Science and Technology Board
Commission on Physical Sciences, Mathematics, and Resources
National Research Council

EXECUTIVE SUMMARY

The United States would benefit significantly from the creation of a national research network (NRN). The National Research Network Review Committee has evaluated the proposal for an NRN submitted by the Office of Science and Technology Policy (OSTP), and the committee strongly supports the concept presented. An NRN would create a computer network infrastructure to provide much-needed support to the scientific research community. Data obtained by the committee regarding current and anticipated research activities demonstrate that an NRN could dramatically improve the productivity and quality of output of the U.S. research community. Through these direct benefits, plus commercial spinoffs from associated computer and network research, an NRN could greatly promote U.S. competitiveness in a multiplicity of disciplines.

The OSTP proposals for an NRN have three phases, a structure the committee supports. In the near term (late 1980s), it would interconnect the existing networks that are fragmentary, overloaded, and poorly functioning. In the second phase (in the early 1990s), it would extend computer network support to researchers at more institutions than are currently served, in particular the smaller institutions not now served well, and it would significantly enhance the range and quality of network-based services provided to researchers. In the final phase (late 1990s) advanced technology and techniques would be deployed to provide a dramatic increase in speed and function to match the power of computers in use at that time.

The information-processing capability on the desk of most researchers today is awesome. The researcher's workstation is an indispensable tool in his daily research endeavors. However, his productivity is being severely hampered by his inability to effectively link his high-performance workstation to the supercomputers, to the large scientific databases, and to his colleagues in a fashion that allows him to incorporate into his research these other resources. What is needed is an accessible, user-friendly, high-speed NRN to provide these connections. The science research community currently sees a serious mismatch between today's high-performance processing engines and low-performance computer

networks. This situation is extremely unfortunate for the United States, the nation that spearheaded networking technology 20 years ago.

National research network users (the scientific research community) have two principal sets of concerns. First, they want--and the committee believes they should receive--an operational network that is stable and predictable, offers high-quality performance, and does not require advanced knowledge of networking technology to use it routinely or to cope with occasional problems. While the OSTP proposals emphasize goals for high speed, the committee equally ranks the concerns for widespread access, user-friendly service-orientation, quality, and performance. Second, users--especially researchers in disciplines that have been less computer- and network-intensive to date--are concerned that federal funding for a research network not diminish federal funding for the conduct of research. The committee shares this concern.

To encourage the development and use of a widely accessible network service, the committee suggests that the NRN might provide (1) a universal basic service at some low to moderate speed for electronic mail at a low cost to users and (2) higher levels of service with an appropriate charging structure. A low charge for basic service would lessen the initial resistance that might otherwise be felt by first-time users and yet discourage the overloading that occurs with free networks.

There will always be a tradeoff between science research users' needs for a stable operational network and network researchers' desires for high performance using the latest networking technology. This latter desire arises both from the developers of network technology, who need a testbed for new ideas, and from that sophisticated class of users with advanced requirements. Because network technology is evolving, there will always be a current networking practice, developing network technologies, and research targets. The challenge underscored by the committee is to manage the progressive introduction of increasingly sophisticated technology without compromising users' ability to conduct the research for which the network is a tool. This may not be possible without substantial investment, and, even with it, compromises may be necessary.

The committee viewed the OSTP proposals as providing a general direction for the NRN, not a set of design specifications (which they were not) to be evaluated. Four areas of concern were addressed: technical considerations, funding, management, and commercial and specialty networks. The recommendations are summarized below:

Technical considerations. The technical demands for the first two phases identified by OSTP are not leading edge items. Indeed, during these phases the NRN will be catching up with the state-of-the-art technology. Phase 1 requires coordination; phase 2 requires careful attention to network management. However, achieving the high-performance network carrying the high volumes of traffic envisioned in phase 3 for the late 1990s requires computer and communications research using a clean-sheet approach that must begin now. The phase 3 network is dominated by technical issues, in particular those associated with network control but also including appropriate technology for switching, routing, multiplexing, processor interfaces, protocols, connection-oriented communications, and layered architectures. All of these areas offer opportunities for technology transfer to industry from the necessary research program.

Funding. There are a number of critical challenges associated with funding for an NRN. These challenges must be recognized and addressed up front, and the transition to a new research funding regime must absolutely lead to a stable funding mechanism or the NRN will never achieve the impact of which it is capable. There are two coupled issues here: (1) funding for network technology and (2) charging for network use. The committee recommends that economic incentives be provided to encourage industry to provide existing network technology to the NRN, but that the government provide funding for advanced network development in cooperation with industrial research efforts where possible. These funds for advanced development should be kept separate from funds allocated to running the operational network. A mechanism for providing (earmarked) funds to users to defray user charges for network use must be established. This mechanism should allow the user to choose among network services. An example of such a mechanism is the use of a voucher for network services.

This transition to user charges requires sensitivity to research users' concerns. For example, one challenge is to introduce charges for NRN use to researchers who may be accustomed to "free" service from such networks as Arpanet. User charges would allocate resources and inhibit excessive demand for use of finite network facilities. A charging structure would also make possible an eventual transition of the service to the private sector. While the economic principles are obvious, the committee recognizes that introducing charges will be difficult and will require compensatory changes in federal grant practices, such as providing for (additional) computer networking costs in research grants.

Management. The scope and scale of the NRN create an imperative for careful attention to key management issues. Good management is absolutely essential for maximizing the effectiveness of the investment in an NRN. It is a critical issue that requires early (and continuing) attention and action, and it is a nontrivial element of the total cost of providing network service. The overall problems are to engineer facilities and services to match demand and tune them to the expressed needs of users, to maintain them in an operational state, and to operate them economically. Network facilities management, operational service management, and process management (for example, introduction and phasing-in of user network services and facilities) each require commitment and resources now. The urgency of the management imperative is underscored by the committee, which points to inadequate service in existing research networks (notably Arpanet and the preexpansion NSFNET) and to the challenges implied by extending access to what could eventually be almost a half-million users in the research community. Poor management and the absence of user services are potential barriers that could prevent the NRN from realizing its potential. To overcome these barriers, an operating structure that does not now exist must be created. It should consist of a funding mechanism through normal government channels, an oversight committee, and an experienced research executive from the private sector for day-to-day management; in addition, one or more user groups should be established to provide input from and output to the user community.

Commercial and Specialty Networks. Discussions of an NRN can easily assume a stand-alone network, but the committee concluded that the proposed network must build on existing and anticipated commercial network facilities and network-based services. It encourages those responsible for developing and implementing the NRN to work with the common carriers and computer vendors; it also encourages long-range planning that would allow for the eventual transition of at least portions of a government-developed NRN to industry, because meeting the service and access goals of the network implies resources beyond the apparent capabilities of federal mission agencies or the National Science Foundation. The commercial network suppliers have decades of experience in network research and in providing accessible and reliable service to millions of users; it would be prudent to leverage this experience for the NRN.

We have a special opportunity with the NRN. We can give an enormous boost to U.S. competitiveness by providing this networking infrastructure. But we must not hesitate, for our foreign competitors are busy at work with their own national research networks (e.g., the RACE program in Europe and national network projects in Canada, Japan, and in newly industrializing countries such as Singapore, and so on). If we are successful in the NRN development, our competitive advantage will continue to grow as this unique infrastructure dramatically improves the quality of output and productivity of the research community of this nation.

U.S. SUPERCOMPUTER VULNERABILITY
BY THE
IEEE SCIENTIFIC SUPERCOMPUTER SUBCOMMITTEE [1]
COMMITTEE ON COMMUNICATIONS AND INFORMATION POLICY

INTRODUCTION

The U.S. created the supercomputer industry. Today U.S. firms jointly are leaders in world supercomputer markets. Nonetheless, U.S. supercomputer firms are vulnerable to a focused strategy by the Japanese targeted on their industry. As documented below, the reasons for the vulnerability of the U.S. firms are complex. To overcome their vulnerability will require a systems solution: an integrated, cooperative effort among industry, universities and government. It appears that such a solution will require coordination by government to a degree seldom if ever achieved in this country in peace time. But the threat is real, and it is not limited to supercomputers. Supercomputers appear to represent only a next step in an on-going process that is gaining momentum. As one can see from the "Visions" for the future promulgated by the Ministry of International Trade and Industry (MITI), Japan seeks a dominant position in the information industries--the computer and communications industries of the world. And it is well on the way to achieving that position.

Many doubt that the U.S. government could play an effective, leadership, policy role in relation to civilian technology/industry problems of this kind. Yet the Japanese government clearly does play such a role in its economy. And Japan is a market-oriented, capitalistic, democracy. By using its approach Japan is overtaking, and often surpassing, the U.S. in field after field. How long can we afford to wait before we respond?

Perhaps the most important question for the U.S. to ponder is: What is the alternative to effective government leadership and support in vital technology/industry matters such as the ones discussed here? Current approaches are not working. The discussion that follows is focused on the supercomputer issue as compounded by the semiconductor issue. Yet, these two issues appear to be only special cases of the broader, generic problem.

---

[1] The IEEE Scientific Supercomputer Subcommittee is a technology-policy subcommittee of the Committee on Communications and Information Policy under the United States Activities Board of the Institute of Electrical and Electronics Engineers, 1111 19th Street, NW, Washington, DC 20036. The members of the Subcommittee are:

Sidney Fernbach, Chairman
Lara Baker
Vito Bongiorno
Alfred E. Brenner
James F. Decker
Duncan Lawrie
Alan McAdams
Kenneth W. Neves
John Ranelletti
John P. Riganati
Stewart Saphier
Paul B. Schneck
Lloyd Thorndyke
Kenneth F. Tiede
Hugh Walsh

John M. Richardson, Chairman, Committee
on Communications and Information Policy

ANALYSIS

Cray and Control Data Corporation (CDC) with its ETA subsidiary are the two long-term U.S. manufacturers of supercomputers. Cray has the largest world market share. It is a stand-alone company, solely dependent on the supercomputer business. CDC and ETA are emerging from financial difficulties and, while more diversified, are not yet nearly as strong as Cray in the supercomputer field. Both Cray and ETA are in a vulnerable position.

Several important factors contribute to the vulnerability of the U.S. supercomputer firms: (1) the current situation in the semiconductor industry and the recent history of US-Japan trade problems in the industry; (2) the rapid strides the Japanese supercomputer firms have made; (3) the fact that Cray and ETA rely on high performance memory devices that now are available only from Japan; (4) the implications of the recent supercomputer trade agreement with Japan; and (5) the fact that U.S. firms are losing a primary market advantage, their differential access to the huge reservoir of applications software that currently is optimized primarily for their installed base of supercomputers.

These factors are discussed in relation to their specific impact on U.S. supercomputer manufacturers; but the process is generic Japanese. As stated in SCIENCE in November, 1985:[2]

> "...In this process, the Japanese begin by picking a few key, high-volume components over which to compete. Then, as the U.S. producers retreat under extreme price-cutting assaults, the Japanese companies extend the fight step-by-step to other items until they have swept competitors out of the most profitable areas. (Emphasis added.)
>
> This has happened in other industries, and in semiconductors the Japanese have all but taken over the field of memory chips..."

Conditions are ripe for supercomputers to represent the next step in the process.

RECENT SEMICONDUCTOR HISTORY

The U.S. created the semiconductor industry. The U.S. invented the VLSI (very large scale integration) Chip. For decades U.S. firms jointly were the leaders in world semiconductor markets. Nonetheless, U.S. semiconductor firms were vulnerable to a focused strategy by the Japanese on DRAMs, dynamic random access memories.

The Japanese developed and manufactured excellent, highly reliable, high quality DRAMs. Then, (as adjudicated by the U.S. Foreign Trade Commission) for more than two years the Japanese dumped these products (sold them below cost) in the U.S. and third markets. During the time that the Japanese were flooding the market with DRAMs, U.S. firms lost in excess of two billion dollars. U.S. merchant manufacturers of semiconductors are stand-alone companies. Effectively they were forced to withdraw from domestic production of high-performance DRAMs.

---

[2] Eliot Marshall, "Fallout from the Trade War in Chips," SCIENCE, November 22, 1985, p. 918.

In other product lines, SRAMs, static random access memories; and EPROMs, erasable programmable read only memory devices; U.S. firms also experienced duress. The U.S. firms responded with legal (as opposed to technological) actions that led to the U.S. Government's imposing tariffs on the Japanese as sanctions. This government effort, while well intended, raised the prices of chips and thus resulted in increased costs to U.S. supercomputer manufacturers and thus has tended to reduce their competitiveness.

More recently, the problem of dumping by the Japanese has been replaced by the current problem of severe shortages of DRAMs in the U.S. These shortages have been shown by experts from the Brookings Institution to result from reduced actual levels of Japanese output that "happen" to accord almost perfectly with the "forecasts" of production levels that had been made by MITI several months before. DRAM prices in the U.S. have sky rocketed--sometimes tenfold, and more. The DRAM business is now enormously profitable to the Japanese producers of semiconductors. This is a classic pattern for those who achieve effective dominance of a market.

During the period of DRAM shortages, exchange rates have been highly favorable to U.S. domestic production. Furthermore, sales of DRAMs in the U.S. market have been highly profitable to suppliers. Nonetheless, (at the time of this writing) U.S. firms were not re-entering the market, either by investing in new plants or even by reopening "mothballed" DRAM plants. Perhaps they recall their recent huge losses too vividly, or perhaps the rapid "decay" of know-how and expertise in so demanding a business explains their lack of response. In any case, a market that is difficult to re-enter is the type of market in which the classic pattern is most successful.

The problems of U.S. firms are rooted in the new-found power of Japanese firms in important aspects of the semiconductor business. In presentations to the authoring IEEE groups, expert observers have stated variously that Japan has "won the silicon war," (August, 1985); and that, "We (the U.S.) failed. We had it (semiconductor leadership) and we lost it!" (May, 1988). Despite these pessimistic assessments and the current dire situation, we believe that this need not (yet) be the case. But the hour is late, and strong action will be required to prevent it.

JAPANESE SUPERCOMPUTER COMPETITION

Cray and CDC/ETA face the three Japanese manufacturers of supercomputers shown on Table I. Not coincidentally, these three firms are also among the leading Japanese manufacturers of semiconductors. Cray and ETA are tiny in comparison to their Japanese rivals--they are small even in relation to the leading U.S. merchant semiconductor manufacturers. (These firms are also shown on Table I.) Given the experience of the semiconductor industry, plus the facts cited in the next paragraphs, it is not hard to understand why the U.S. supercomputer firms are vulnerable to the "next step" by the Japanese.

The sum of the current annual supercomputer outputs of the three Japanese giants approximates the number, though not yet the dollar value, of Cray's output, and greatly exceeds that of CDC/ETA. As Table I demonstrates, the Japanese firms are not stand-alone manufacturers of supercomputers. Their operations are integrated over multiple production stages and across several product lines.

Table II shows the number of supercomputers of each manufacturer installed and on order as of late 1987. It is important to note that both the CDC 205 and the Cray 1 are now obsolete, while the ETA 10's and Cray Y-MPs (not listed) are just becoming operational. Table III shows the peak megaflops per processor for each manufacturer as well as the megaflops per total (multiprocessor) system.

Table I

Market Segments Participated in By Selected U.S. And Japanese Manufacturers of Supercomputers and/or Semiconductors

	CRAY	CDC/ETA	TI INTEL Motorola AMD ETC.	AT&T	IBM	NEC	FUJITSU	HITACHI
Super-computers	x	x			RE*	x	x	x
Main Frame Computers		x			x	x	x	x
Intermediate Computers					x	x	x	x
Mini Computers		x		x	x	x	x	x
Micro Computers					x	x	x	
Consumer Electronics						x	x	x
Semiconductors								
- Merchant			x			x	x	x
- Captive				x	x	x	x	x

---

* Re-entering.

Table II
Supercomputer Installations

System	No. Installed	No. on Order	FCS
CDC/ETA			
CDC Cyber 205	37	0	1980
ETA 10 series	2	8	1987
Cray			
Cray-2	8	4	1985
X-MP series	117	43	1984
Cray-1	58	0	1976
Fujitsu			
VP series	56	5	1984
Hitachi			
S-810 & 820	19	4	1983
NEC			
SX2 series	12	7	1985

Source: Adapted from The Gartner Group, November, 1987.

Table III
Peak Performance Rates

COMPUTER	SINGLE CPU PEAK 64-BIT MFLOP RATE	ALLOWING MULTIPLE CPUs
CRAY-1	160	160
CRAY X-MP	233	932 (4-CPUs)
CRAY-2	488	1,952 (4-CPUs)
CRAY-3 (1989)	1,000	16,000 (16-CPUs) est.
CYBER 205	200 (2-PIPE)	400 (4-PIPE)
ETA-10 (1986)	350	1,400 (4-CPUs)
ETA-10/E	415	1,660 (4-CPUs)
ETA-10/G (1988)	643	5,142 (8-CPUs)
FUJ. VP 100	271	271
FUJ. VP 200	533	533
FUJ. VP 400	1,067	1,067
HIT. S-810/20	630	630
HIT. S-820/80	2,000	N/A
IBM 3090/VF	116	696 (6-CPUs)
NEC SX-1	570	570
NEC SX-2	1,300	1,300
NEC SX-3 (1989)	5,000	20,000 (4-CPUs)

Source: Adapted from "Computational Fluid Dynamics: Algorithms and Supercomputers," by W. Gentzsch and K. Neves, NATO AGARDograph No. 311, March, 1988.

The facts in these tables are consistent with the following summary. The Japanese manufacturers, especially Fujitsu, have made rapid progress in installing systems during the four short years they have been producing supercomputers (Table II). Their systems are to date single processor systems (Table III). This latter point relates to two elements: (A) the individual processors in the Japanese supercomputers are more powerful today than any U.S. manufacturer currently expects to deliver in less than half a decade; but, (B) the Japanese manufacturers today lag somewhat in adapting to parallel processing techniques (though they appear to lead in the complementary field of vectorization and vector processing; see below).

These facts imply that multiprocessor versions of the Japanese supercomputers, especially the newest Hitachi and NEC (Nippon Electric Corporation) machines, could far outstrip the best of the U.S. machines once the Japanese catch up in parallel processing, an area where they are hard at work. NEC has announced parallelism for its SX-2 series and is said to have completed the circuit design for its 20 gigaflop, four processor SX-3 series currently in development.

In a recent special issue (March 3, 1988) of the trade journal Electronics entitled "Inside Technology,"[3] these facts are reiterated in very strong language:

> "...The Japanese stick to simpler single-processor architectures. And their supercomputers are the speediest one-processor machines in the world--by far--thanks primarily to advanced semiconductor technology. The only way U.S. supercomputers can come near their processing rates is in multiprocessor configurations."

The following quotes also are in the article in a box entitled, "Where Japan Shines: Vectorizing Compilers."[4]

> "...Japanese...compilers make it relatively easy for a user with his own Fortran program to wring near maximum performance from a supercomputer's vector processor without resorting to special vector-processor programming methods.
>
> Automatic vectorization has been a Japanese strong suit for some time, and it is made easier by the Japanese's use of single-processor architectures in supercomputers. Compiler writers need not face the difficulties of automatic parallelization."

The latter quote points up an additional area of Japanese strength; the systems software of the Japanese supercomputers is already world class. NEC's Fortran compiler is reputed to be the best in the world. If that is in fact the case, it would have surpassed the previous best--that from Hitachi.

---

[3] Charles L. Cohen, "Japan Focuses on Simple but Fast Single-Processor Supercomputers," Electronics, March 3, 1988, p. 57.

[4] Ibid. p. 58.

U.S. SUPERCOMPUTER FIRMS RELIANCE ON JAPANESE SEMICONDUCTORS

Already the aggressive pricing strategies of the Japanese in supercomputers (discounts of up to 80% to universities and others) combined with inversely aggressive semiconductor strategies (refusal to export their latest high performance component devices), threaten the existence of Cray and ETA. The highest performance memory and bipolar logic components useful for supercomputers are no longer manufactured in the U.S.; they are available only from Japan. The managements of Cray and ETA have been quoted in the press at various times as stating that these Japanese components are, "not yet available for export" from Japan to Cray or ETA as devices--but they are available to end users in the Japanese supercomputer systems. Those systems are definitely available for export.

This continues the familiar, oft-repeated pattern. Japanese firms plan and act in accord with long-range strategic goals. When they achieve a technological advantage in one area, they use that advantage to ensure their advance into new areas. They have targeted supercomputers as the next high tech area in which to establish a dominant position, and they have made enormous strides toward that goal as Tables II and III demonstrate.

THE IMPLICATIONS OF THE SUPERCOMPUTER TRADE AGREEMENT WITH JAPAN

The supercomputer trade agreement with Japan contains the seeds which can further undermine the position of U.S. manufacturers. It requires of Japan a number of actions which, if made reciprocal to the U.S., could facilitate the entry of Japanese supercomputers into the U.S. market.

The reasoning that supports this conclusion proceeds through the following steps. Almost by definition, no one knows how to use a new generation of supercomputers efficiently. New systems generally embody the latest technology and architectural design in order to achieve their state-of-the-art computing speeds. In this country, the national laboratories historically have been "partners" in supercomputer development. The laboratories have developed the applications and forced the evolution of the operating systems that together permit the machines to work effectively. In other words, the laboratories' expert users have assisted greatly in facilitating improved productivity of the systems as these users learned more about the characteristics of the new systems and how to make optimal use of them. Now with the new NSF programs that have made supercomputers available to universities, universities, too, are becoming partners in these efforts. Often the U.S. manufacturer has recognized the "partnership contribution" of these U.S. institutions through price concessions, "buybacks" of time or other means. The same has been true in Japan.

The trade agreement requires the Japanese in procurements of supercomputers by their government or their universities to give equal preference to U.S. manufactured supercomputers. In essence, the Japanese have agreed to avoid "unfair" pricing on the part of the Japanese manufacturers. But what may seem "unfair" to a trade administrator looks to a manufacturer like "recognition of partnership" with a government agency or university.

While negotiating this agreement with one hand, the U.S. government has blocked the sale of a NEC supercomputer to M.I.T. with the other. To date Japanese supercomputers have been virtually shut out of U.S. government and university markets. It appears that government agencies, especially the Department of Defense, intend to keep it that way.

How long can such a pattern last? Is it realistic to believe that the Japanese can long be required to give unbiased consideration to U.S. supercomputers while Japanese machines are foreclosed from U.S. institutions? If the requirements we now impose on the Japanese were made reciprocal for U.S. universities and government laboratories, Japanese supercomputers would have to be acceptable to those agencies on non-discriminatory terms. The Japanese systems are already a match for the U.S. systems and soon could surpass them in performance. Price thus is likely to be a major determinant of choice; the Japanese manufacturers have been more than willing to sacrifice short-term profit for long-term market penetration.

The new requirements could greatly facilitate the entry of the Japanese into the U.S. At the same time such requirements could disrupt the implicit partnership between U.S. manufacturers and U.S. government laboratories and/or U.S. universities, and transfer the benefits of partnership to Japanese firms. In response to a low bid, U.S. National Laboratories could be required to become partners to Japanese firms in perfecting their systems for penetration of U.S. markets.

INCREASING "PORTABILITY" OF SOFTWARE

Paradoxically, desirable developments long sought by users on another front are hastening the progress of the Japanese. It is becoming increasingly easy to move applications from the machine of one supercomputer vendor to that of another. There are four elements which are facilitating software "portability:"

(1) Standardized versions of FORTRAN;
(2) Improved FORTRAN optimizing compilers;
(3) Broad-spectrum libraries of algorithms optimized for each vendor's systems;
(4) A common operating system--UNIX--now being implemented on supercomputers.

While in the past each supercomputer vendor had a "captive audience" of users, now it is increasingly possible to move applications between and among vendors. As the industry leader with the largest installed base, Cray had benefited from the prior situation in the industry; CDC did also, but to a lesser extent. But portable software does provide substantial benefit to users of supercomputers, it greatly diminishes this advantage previously enjoyed by the U.S. firms.

New entry into the field of supercomputing also is facilitated by increased applications portability. As the most recent entrants, the Japanese vendors are the greatest beneficiaries of these changes.

Cray itself also benefits from portability since portability helps Cray provide compatibility to its own multiple product lines. Use of UNIX permits Cray to focus on a single operating system thus reducing its development costs. The most significant effects, however, are those on the new entrants.

A further benefit to users is also achieved. Through UNIX, an operating environment uniform from work station to minicomputer, to minisupercomputer, to supercomputer would become possible. The Los Alamos National Laboratory is one of the organizations working hard to bring this about. One of its major objectives is to bring the myriad, powerful software innovations pioneered in the highly competitive fields of commercial computing over into the more rarified arena of supercomputing--with benefit to U.S. manufacturers as well as to users. The other side of the coin, of course, is that these benefits accrue to all vendors of supercomputers--especially to the Japanese.

THE IMPLICATIONS OF THE PHENOMENA OF SEMICONDUCTOR DISADVANTAGE AND THE PRESENCE OF PORTABILITY

At this point two threads come together. Once general portability of applications is achieved, the avenue to continued market leadership by U.S. firms would be solely through technological leadership. But how can U.S. firms simultaneously rely on componentry manufactured by their competitors, the Japanese, and assure their customers that they, the U.S. firms, can maintain technological leadership? The Japanese dominate completely in semiconductor memory (and in bi-polar logic). The American firms have each sought logic components of usual technology to advance their competitive positions: Cray with Gallium Arsenide chips for the Cray III, ETA with ultra-dense C-MOS chips cooled with liquid nitrogen for the ETA-10. (The latter chips are available from Honeywell as a result of that firm's participation in the VHSIC "Very High Speed Integrated Circuit" program of the Department of Defense. To date, these have not been manufactured to the specifications required for ETA to meet its advertised cycle-time of 7 nanoseconds.) But we see also that despite these efforts, at their rated performance levels the U.S. systems lose the megaflops/processor race (Table III). And each firm's effort is based on a favorable "spike" in an otherwise dismal U.S. semiconductor picture--which as yet shows no signs of improvement sufficient to reverse its overall, rapid decline.

From experience we can see that once the Japanese fully consolidate their semiconductor technology leadership, in time U.S. firms would be most fortunate if they found themselves to be only a generation or so behind them. The experience of Asia's "Four Tigers," Korea, Taiwan, Hong Kong, and Singapore confirms this point. The Japanese have consistently resisted the requests of the tigers for newer Japanese technologies. The tigers are explicitly told by the Japanese that this is because Japan looks upon them as "potential future competitors." As a rule, the Japanese refuse to license them for any technology less than five years old.

LONGER RUN IMPLICATIONS

The "vector facility" which provides vector capability to IBM's 3090 Mainframes, today blurs the boundary between the mainframe and the supercomputer. In any case, IBM has demonstrated its renewed interest in the field of supercomputing. This confirms the field's growing commercial importance--but IBM is not yet well established as a vendor of supercomputers.

The Japanese are established in supercomputing. They are challenging IBM in mainframes: Hitachi, with its plug compatible CPU's marketed by National Semiconductor's subsidiary, NAS; Fujitsu/Amdahl with their CPU's marketed through Amdahl Inc. (which recently experienced its best year). NEC's ACOS series is marketed in the U.S. through Honeywell. The mainframes from all three of these vendors also have some vector capability. Most Japanese supercomputers are "IBM-compatible." This is especially true of those from Fujitsu and Hitachi. (And the Fortran for NEC's SX-2 series is known as its "IBM version." For the SX-3 series NEC plans both an "IBM version" and a "Cray version.") Mainframes are growing larger and faster. But the Japanese appear to have intercepted that market from above through their supercomputers.

CONCLUSION

Once they were to achieve preeminence in supercomputers, the Japanese would have established themselves as the manufacturers of the most powerful computer systems, with supercomputers installed with the "Gold Chip" accounts of the world. They could then use that base and prestige to cement their move into the more lucrative market for commercial data processing systems. That is a long-term goal.

The threat is not just to supercomputers. The threat is to the entire U.S. base of computer systems. It is ironic that such cascading impacts could originate from dominance of a product area--DRAMs--that has reached almost commodity status. Yet such impacts appear not only possible, but likely. There is no denying the enormous impact that current DRAM shortages are already having on U.S. systems manufacturers. Their component costs are rising sharply. Many are unable to meet their orders due solely to lack of DRAMs.

The most recent developments with Sematech further confirm the depth of the crisis in U.S. semiconductor manufacturing at two levels. First, IBM and AT&T each has agreed to donate an advanced technology to Sematech. For IBM it is the four megabit DRAM. For AT&T it is the technologically equivalent one megabit SRAM. That these two arch rivals would see their way clear to share technology with each other and with others in the U.S. semiconductor industry demonstrates the depth of their concern. In part this is because these integrated firms buy their equipment for manufacturing semiconductors from others. And the Japanese are dangerously close to becoming the sole source of that equipment. Second, it is not encouraging that (as of this writing) this fledgling effort at cooperation by the semiconductor industry, Sematech, remains understaffed, organizationally in disarray, and well behind its self-imposed schedule.

At another level, it must be recognized that there is no assurance that even a Sematech highly successful in relation to its stated goals was set up to be coordinated with needed actions in relation to supercomputers--or with any other elements required for an effective national technology policy. This brings us back to the point made in the opening paragraph.

REQUIRED: A SYSTEM SOLUTION

The solution to the problems of the U.S. supercomputer manufacturers lies not in imploring the Japanese not to pursue their advantage, but in the U.S. taking positive actions at home to insure that the Japanese don't succeed at our expense. The U.S. must get its technology base in order: in semiconductors, in supercomputers and in software. No current initiative--including Sematech or IBM's own recent focus on the field of supercomputers--assures that these objectives will be achieved. A system solution cannot be achieved through discrete, uncoordinated initiatives, no matter how worthy.

Several studies over recent years make it clear that a solution to the problems of U.S. civilian technology/industry/trade policy, such as the matters discussed in this paper, will require expert coordination and leadership.[5] It seems clear that to be effective, these can only by provided at the national level. Only at this level would it be possible for a true systems approach to be taken to the problems facing this country in the international competitiveness of its industry and of its high technology base. Action is required on long term R&D policies: it is necessary to integrate the activities of government, industry and universities; it is necessary to achieve timely, effective technology transfer from stage to stage in the process; it's necessary to ensure careful attention to the provision of suitable people-power to each stage in the process, an issue involving not just numbers, but also appropriate intellectual capital.

---

[5] For example, see "Global Competition: The New Reality," the Report of the President's Commission on Industrial Competitiveness, January, 1985; Superintendent of Documents, Washington, DC, 20402; and "The Technological Dimensions of International Competitiveness," May, 1988, Office of Administration, Finance, and Public Awareness, National Academy of Engineering, 2101 Constitution Ave, NW, Washington, DC 20418.

The challenge is to find an acceptable institutional framework in which government, industry and academia can pursue these objectives to the long-run benefit of the nation as a whole. A necessary requirement of such a framework is: it must ensure that economic and technological decisions be taken on the basis of economic and technological--as opposed to mainly political or military--criteria. The answer may well require that the coordination and leadership functions be vested in a new, lean, expert, civilian agency of government that is capable of focusing on the longer term national interest. Only through a coordinated approach to all of the above issues will we be able to ensure a strong U.S. base for innovation, productivity, and international competitiveness--a base in which supercomputers constitute a vital factor.

○